



Epsilon Programmer's Editor User's Manual and Reference

Version 14.00 - Reference Edition

This is revision 14.00a of the manual.

It describes version 14.00 of Epsilon and EEL.

Copyright © 1984, 2020 by Lugaru Software Ltd.

All rights reserved.

Lugaru Software Ltd.

1645 Shady Avenue

Pittsburgh, PA 15217

TEL: (412) 421-5911

E-mail: support@lugaru.com or sales@lugaru.com

LIMITED WARRANTY

THERE ARE NO WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE FOR EITHER THE INSTRUCTION MANUAL, OR FOR THE EPSILON PROGRAMMER'S EDITOR AND THE EEL SOFTWARE (COLLECTIVELY, THE "SOFTWARE").

Lugaru warrants the medium on which the Software is furnished to be free from defects in material under normal use for ninety (90) days from the original date of purchase, provided that the limited warranty has been registered by mailing in the registration form accompanying the Software.

LIMITED LIABILITY AND RETURN POLICY

Lugaru will be liable only for the replacement of defective media, as warranted above, which are returned shipping prepaid to Lugaru within the warranty period. Because Lugaru cannot anticipate the intended use to which its Software may be applied, it does not warrant the performance of the Software. LUGARU WILL NOT BE LIABLE FOR ANY SPECIAL, INDIRECT, CONSEQUENTIAL OR OTHER DAMAGES WHATSOEVER. However, Lugaru wants you to be completely satisfied with the Software. Therefore, THE ORIGINAL PURCHASER OF THIS SOFTWARE MAY RETURN IT UNCONDITIONALLY TO LUGARU FOR A FULL REFUND FOR ANY REASON WITHIN SIXTY DAYS OF PURCHASE, PROVIDED THAT THE PRODUCT WAS PURCHASED DIRECTLY FROM LUGARU SOFTWARE LTD.

COPYRIGHT NOTICE

Copyright © 1984, 2020 by Lugaru Software Ltd. All rights reserved.

Lugaru Software Ltd. recognizes that users of Epsilon may wish to alter the EEL implementations of various editor commands and circulate their changes to other users of Epsilon. Limited permission is hereby granted to reproduce and modify the EEL source code to the commands provided that the resulting code is used only in conjunction with Lugaru products and that this notice is retained in any such reproduction or modification.

TRADEMARKS

"Lugaru" and "EEL" are trademarks of Lugaru Software, Ltd. "Epsilon" is a registered trademark of Epsilon Data Management, Inc. Lugaru Software Ltd. is licensed by Epsilon Data Management, Inc. to use the "Epsilon" mark in connection with computer programming software. There is no other affiliation or association between Epsilon Data Management, Inc. and Lugaru Software Ltd. "Brief" is a registered trademark of Borland International.

SUBMISSIONS

Lugaru Software Ltd. encourages the submission of comments and suggestions concerning its products. All suggestions will be given serious technical consideration. By submitting material to Lugaru, you are granting Lugaru the right, at its own discretion and without liability to you, to make any use of the material it deems appropriate.

Note to Our Users

Individual copies of Epsilon aren't protected with a formal license agreement, but by copyright law. In addition to the copying for backup permitted under copyright law, Lugaru grants you, the end-user, certain other rights, as explained on this page.

It describes the rules for installing a single purchased copy of Epsilon on multiple computers, and related matters. These rules apply to all copies of Epsilon purchased by an end-user and not subject to a written license agreement.

Each copy of Epsilon includes packages for various operating systems or distributions, such as a Windows package, a Debian Linux package, and a Macintosh package.

You may install a single purchased copy of Epsilon on up to four computers under your control, either installing the same package on each, or a different package on each, or any combination, as long as you're the only one using any of these packages. Two individuals may not share a single copy of Epsilon if there is any chance both individuals might use that copy of Epsilon at the same time, even by using separate packages on separate computers.

You may not split a single purchased copy of Epsilon into its separate packages and sell them separately.

If you purchase an update to Epsilon, it becomes part of the same copy. You may not (for example) buy Epsilon 10, update to Epsilon 11, and then sell Epsilon 10 while retaining Epsilon 11. The update does not count as a separate copy and must accompany the original version if sold.

We hope that you will respect our efforts, and the law, and not allow illegal copying of Epsilon.

We wish to thank all of our users who have made Epsilon successful, and extend our welcome to all new users.

Steven Doerfler
Lugaru Software, Ltd.

*We produced this manual using the Epsilon Programmer's Editor and the T_EX typesetting system.
Duane Bibby did the illustrations.*

Contents

1	Welcome	1
1.1	Introduction	1
1.2	Features	1
2	Getting Started	5
2.1	Installing Epsilon for Windows	5
2.2	Installing Epsilon for Unix	5
2.3	Installing Epsilon for macOS	6
2.3.1	Epsilon for macOS Keyboard Issues	7
2.3.2	Epsilon for macOS Startup	7
2.3.3	Epsilon for macOS and Privacy-protected Folders	8
2.4	Installing Epsilon for DOS or Epsilon for OS/2	9
2.5	Tutorial	9
2.6	Invoking Epsilon	9
2.7	Configuration Variables: The Environment and The Registry	10
2.7.1	How Epsilon Finds its Files	12
2.7.2	The Customization Directory	13
2.8	Epsilon Command Line Flags	13
2.9	File Inventory	17
3	General Concepts	23
3.1	Buffers	23
3.2	Windows	23
3.3	Epsilon's Screen Layout	23
3.4	Different Keys for Different Uses: Modes	25
3.5	Keystrokes and Commands: Bindings	26
3.6	Repeating: Numeric Arguments	27
3.7	Viewing Lists	27
3.8	Typing Less: Completion & Defaults	28
3.9	Command History	30
3.10	Mouse Support	31
3.11	The Menu Bar	32
3.11.1	Customizing Epsilon's Menu	33
4	Commands by Topic	37
4.1	Getting Help	37
4.1.1	Info Mode	39
4.1.2	Web-based Epsilon Documentation	41
4.2	Moving Around	42

4.2.1	Simple Movement Commands	42
4.2.2	Moving in Larger Units	42
4.2.3	Searching	45
4.2.4	Bookmarks	50
4.2.5	Tags	50
4.2.6	Source Code Browsing Interface	52
4.2.7	Comparing Two Buffers	54
4.2.8	Comparing Many Files	56
4.3	Changing Text	57
4.3.1	Inserting and Deleting	57
4.3.2	The Region, the Mark, and Killing	59
4.3.3	Clipboard Access	61
4.3.4	Rectangle Commands	62
4.3.5	Capitalization	63
4.3.6	Replacing	64
4.3.7	Regular Expressions	66
4.3.8	Rearranging	75
4.3.9	Indenting Commands	78
4.3.10	Aligning	80
4.3.11	Automatically Generated Text	81
4.3.12	Spell Checking	81
4.3.13	Hex Mode	83
4.4	Language Modes	84
4.4.1	Asm Mode	85
4.4.2	Batch Mode	85
4.4.3	C Mode	86
4.4.4	Configuration File Mode	90
4.4.5	GAMS Mode	90
4.4.6	HTML, XML, and CSS Modes	91
4.4.7	Ini File Mode	93
4.4.8	Makefile Mode	93
4.4.9	NSIS Mode	93
4.4.10	Perl Mode	94
4.4.11	PHP Mode	95
4.4.12	PostScript Mode	95
4.4.13	Python Mode	95
4.4.14	R Mode	96
4.4.15	Shell Mode	96
4.4.16	Tcl Mode	97
4.4.17	TeX and LaTeX Modes	97
4.4.18	VHDL Mode	100
4.4.19	Visual Basic Mode	100
4.5	More Programming Features	100
4.5.1	Navigating in Source Code	101
4.5.2	Pulling Words	101
4.5.3	Accessing Help	102
4.5.4	Context-Sensitive Help	102
4.5.5	Commenting Commands	104
4.6	Fixing Mistakes	105
4.6.1	Undoing	105
4.6.2	Interrupting a Command	106

4.7	The Screen	106
4.7.1	Display Commands	106
4.7.2	Horizontal Scrolling	108
4.7.3	Windows	108
4.7.4	Customizing the Screen	110
4.7.5	Fonts	112
4.7.6	Setting Colors	113
4.7.7	Code Coloring	114
4.7.8	Window Borders	115
4.7.9	The Bell	116
4.8	Buffers and Files	117
4.8.1	Buffers	117
4.8.2	Files	118
4.8.3	File Variables	127
4.8.4	Internet Support	130
4.8.5	Unicode Features	136
4.8.6	Printing	137
4.8.7	Extended file patterns	138
4.8.8	Directory Editing	139
4.8.9	Buffer List Editing	143
4.9	Starting and Stopping Epsilon	144
4.9.1	Session Files	144
4.9.2	File Associations	146
4.9.3	Sending Files to a Prior Instance	146
4.9.4	MS-Windows Integration Features	147
4.10	Running Other Programs	149
4.10.1	The Concurrent Process	150
4.10.2	Compiling From Epsilon	153
4.11	Repeating Commands	155
4.11.1	Repeating a Single Command	155
4.11.2	Keyboard Macros	155
4.12	Simple Customizing	157
4.12.1	Bindings	157
4.12.2	Brief Emulation	159
4.12.3	CUA Keyboard	159
4.12.4	Variables	159
4.12.5	Saving Changes to Bindings and Variables	162
4.12.6	Command Files	164
4.13	Advanced Topics	168
4.13.1	Changing Commands with EEL	168
4.13.2	Updating from an Old Version	170
4.13.3	Keys and their Representation	172
4.13.4	Customizing the Mouse	175
4.14	Miscellaneous	176
5	Alphabetical Command List	179
6	Variables	265
7	Changing Epsilon	387
8	Introduction to EEL	391

8.1	Epsilon Extension Language	391
8.2	EEL Tutorial	391
9	Epsilon Extension Language	397
9.1	EEL Command Line Flags	397
9.2	The EEL Preprocessor	398
9.3	Lexical Rules	401
9.3.1	Identifiers	401
9.3.2	Numeric Constants	401
9.3.3	Character Constants	401
9.3.4	String Constants	402
9.4	Scope of Variables	402
9.5	Data Types	403
9.5.1	Declarations	404
9.5.2	Simple Declarators	405
9.5.3	Pointer Declarators	405
9.5.4	Array Declarators	406
9.5.5	Function Declarators	407
9.5.6	Structure and Union Declarations	407
9.5.7	Complex Declarators	409
9.5.8	Typedefs	409
9.5.9	Type Names	409
9.6	Initialization	410
9.7	Statements	412
9.7.1	Expression Statement	412
9.7.2	If Statement	413
9.7.3	While, Do While, and For Statements	413
9.7.4	Switch, Case, and Default Statements	413
9.7.5	Break and Continue Statements	414
9.7.6	Return Statement	414
9.7.7	Save_var and Save_spot Statements	414
9.7.8	On_exit Statement	415
9.7.9	Goto and Empty Statements	416
9.7.10	Block	416
9.8	Conversions	416
9.9	Operator Grouping	417
9.10	Order of Evaluation	418
9.11	Expressions	418
9.11.1	Constants and Identifiers	418
9.11.2	Unary Operators	419
9.11.3	Simple Binary Operators	420
9.11.4	Assignment Operators	421
9.11.5	Function Calls	422
9.11.6	Miscellaneous Operators	422
9.12	Constant Expressions	423
9.13	Global Definitions	423
9.13.1	Key Tables	424
9.13.2	Color Classes	425
9.13.3	Function Definitions	427
9.14	Differences Between EEL And C	429
9.15	Syntax Summary	430

10 Primitives and EEL Subroutines	439
10.1 Buffer Primitives	439
10.1.1 Changing Buffer Contents	439
10.1.2 Moving Text Between Buffers	441
10.1.3 Getting Text from a Buffer	441
10.1.4 Spots	442
10.1.5 Narrowing	444
10.1.6 Undo	445
10.1.7 Searching Primitives	446
10.1.8 Moving by Lines	451
10.1.9 Other Movement Functions	452
10.1.10 Sorting Primitives	453
10.1.11 Other Formatting Functions	453
10.1.12 Comparing	454
10.1.13 Managing Buffers	456
10.1.14 Catching Buffer Changes	457
10.1.15 Listing Buffers	458
10.2 Display Primitives	459
10.2.1 Creating & Destroying Windows	459
10.2.2 Window Resizing Primitives	461
10.2.3 Preserving Window Arrangements	461
10.2.4 Pop-up Windows	462
10.2.5 Pop-up Window Subroutines	464
10.2.6 Window Attributes	465
10.2.7 Buffer Text in Windows	466
10.2.8 Window Titles and Mode Lines	468
10.2.9 Normal Buffer Display	471
10.2.10 Displaying Status Messages	477
10.2.11 Printf-style Format Strings	480
10.2.12 Other Display Primitives	482
10.2.13 Highlighted Regions	483
10.2.14 Character Coloring	486
10.2.15 Code Coloring Internals	488
10.2.16 Colors	492
10.3 File Primitives	495
10.3.1 Reading Files	495
10.3.2 Writing Files	497
10.3.3 Line Translation	498
10.3.4 Character Encoding Conversions	500
10.3.5 More File Primitives	502
10.3.6 File Properties	505
10.3.7 Low-level File Primitives	507
10.3.8 Directories	508
10.3.9 Manipulating File Names	510
10.3.10 Internet Primitives	514
10.3.11 Tagging Internals	519
10.4 Operating System Primitives	519
10.4.1 System Primitives	519
10.4.2 Window System Primitives	522
10.4.3 Timing	527
10.4.4 Calling DLLs (Windows Only)	528

10.4.5	Running a Process	529
10.5	Control Primitives	534
10.5.1	Control Flow	534
10.5.2	Character Types	537
10.5.3	Examining Strings	539
10.5.4	Modifying Strings	541
10.5.5	Byte Arrays	542
10.5.6	Memory Allocation	542
10.5.7	The Name Table	543
10.5.8	Built-in and User Variables	545
10.5.9	Buffer-specific and Window-specific Variables	548
10.5.10	Bytecode Files	549
10.5.11	Starting and Finishing	550
10.5.12	EEL Debugging and Profiling	552
10.5.13	Help Subroutines	554
10.6	Input Primitives	555
10.6.1	Keys	555
10.6.2	The Mouse	559
10.6.3	Window Events	565
10.6.4	Completion Subroutines	566
10.6.5	Other Input Functions	572
10.6.6	Dialogs	574
10.6.7	The Main Loop	578
10.6.8	Bindings	579
10.7	Defining Language Modes	584
10.7.1	Language-specific Subroutines	589
11	Error Messages	591
A	Index	593

Chapter 1

Welcome



1.1 Introduction

Welcome! We hope you enjoy using Epsilon. We think you'll find that Epsilon provides power and flexibility unmatched by any other editor for a personal computer.

Epsilon has a command set and general philosophy similar to the EMACS-style editors used on many different kinds of computers. If you've used an EMACS-style editor before, you will find Epsilon's most commonly used commands and keys familiar. If you haven't used an EMACS-style editor before, you can use Epsilon's tutorial program. Chapter 2 tells you how to install Epsilon and how to use the tutorial program.

1.2 Features

- Full screen editing with an EMACS-style command set.
- An exceptionally powerful embedded programming language, called EEL, that lets you customize or extend the editor. EEL provides most of the expressive power of the C programming language.
- You can invoke your compiler or “make” program from within Epsilon, then have Epsilon scan the output for error messages, then position you at the offending line in your source file. See page 153.
- An undo command that lets you “take back” your last command, or take back a sequence of commands. The undo facility works on both simple and complicated commands. Epsilon has a redo command as well, so you can even undo your undo's. See page 105.
- Very fast redisplay. We designed Epsilon specifically for the personal computer, so it takes advantage of the high available display bandwidth.
- Epsilon can dynamically syntax-highlight source code files written in many different languages, showing keywords in one color, functions in another, string constants in a third, and so forth.
- Epsilon can finish typing long identifier names for you.
- You can interactively rearrange the keyboard to suit your preferences, and save the layout so that Epsilon uses it the next time. Epsilon can also emulate the Brief text editor's commands, or use a CUA-style keyboard (like various Windows programs).
- You can edit a virtually unlimited number of files simultaneously.
- Epsilon understands Internet URLs and can asynchronously retrieve and send files via FTP. It also includes support for Telnet, SSH, SCP, and various other protocols.
- Epsilon provides a multi-windowed editing environment, so you can view several files simultaneously. You can use as many windows as will fit on the display. See page 108.
- Under Windows, Epsilon provides a customizable tool bar.
- The ability to run other programs from within Epsilon in various ways. See page 149.
- The ability to run some classes of programs concurrently with the output going to a window. Details begin on page 150.
- An extensive on-line help system. You can get help on what any command does, what any key does, and on what the command executing at the moment does. And Epsilon's help system will automatically know about any rearrangement you make to the keyboard. See page 37.

- An extensible “tags” system for many programming languages that remembers the locations of subroutine and variable definitions. You provide a subroutine name, for instance, and Epsilon takes you to the place that defines that subroutine. Alternatively, you can position the cursor on a function call, hit a key, and jump right to the definition of that function. See page 50.
- Completion on file names and command names. Epsilon will help you type the names of files and commands, and display lists of names that match a pattern that you specify. You can complete on many other classes of names too. This saves you a lot of typing. See page 28.
- Support for Unicode files and files using a variety of other character sets.
- Under Windows, you can drag and drop files or directories onto Epsilon’s window, and Epsilon will open them.
- Commands to manipulate words, sentences, paragraphs, and parenthetical expressions. See the commands starting on page 42.
- Indenting and formatting commands. Details start on page 77.
- A kill ring to store text you’ve previously deleted. You can set the number of such items to save. See page 59.
- A convenient *incremental search* command (described on page 45), as well as regular searching commands, and search-and-replace commands.
- Regular expression searches. With regular expressions you can search for complex patterns, using such things as wildcards, character classes, alternation, and repeating. You can even search based on syntax highlighting, finding only matches in a programming language comment or string, or using Unicode property names.
- A fast *grep* command that lets you search across a set of files. See page 48. You can also replace text in a set of files.
- Extended file patterns that let you easily search out files on a disk.
- A directory editing command that lets you navigate among directories, copying, moving, and deleting files as needed. It even works on remote directories via FTP or SCP.
- Fast *sort* commands that let you quickly sort a buffer. See page 75.
- A powerful *keyboard macro* facility (see page 155), that allows you to execute sequences of keystrokes as a unit, and to extend the command set of the editor. You’ll find Epsilon’s keyboard macros very easy to define and use.
- Commands to compare two files and find the differences between them. You can compare character-by-character or line-by-line, displaying results in a variety of formats. See page 54.
- You can choose from a variety of built-in screen layouts, making Epsilon’s screen look like those of other editors, or customize your own look for the editor.

Chapter 2

Getting Started



This chapter tells you how to install Epsilon on your system and explains how to invoke Epsilon. We also describe how to run the tutorial, and list the files in an Epsilon distribution.

2.1 Installing Epsilon for Windows

To install Epsilon for Windows, run the Windows installer you downloaded. The usual version installs Epsilon system-wide and requires admin privileges, but we also include a setup program that does not, and installs Epsilon just for the current user.

Each installation program installs the GUI version of Epsilon for Windows, and the Windows Console version. We named the Windows GUI version `epsilon.exe` and the console version `epsilonc.exe`.

The installation program creates program items to run Epsilon. You can recreate them, set up file associations, change the registration information shown in Epsilon's About box, and do similar reconfiguration tasks by running Epsilon's `configure-epsilon` command.

The installer also sets the registry entry `Software\Lugaru\Epsilon\EpsPathversion` in the `HKEY_CURRENT_USER` hierarchy to include the name of the directory in which you installed Epsilon (where *version* represents Epsilon's version number).

You can uninstall Epsilon by using the "Programs and Features" Control Panel, or from `configure-epsilon`.

2.2 Installing Epsilon for Unix

Epsilon includes a version for Linux and a separate version for FreeBSD. We describe them (and the macOS version) collectively as the "Unix" version of Epsilon. The Linux version is available in `.rpm` and `.deb` formats, while the FreeBSD version is provided in `.tar.bz2` format. See the download page for installation instructions.

Some versions of Epsilon use a helper program to access certain shared library files from the glibc 2.1 NSS subsystem. If necessary, the installation script will compile a helper program to provide Epsilon with these services.

Epsilon runs as an X11 program when run under the X11 windowing system, and as a text program outside of X. Epsilon knows to use X when it inherits a `DISPLAY` environment variable. You can override Epsilon's determination by providing a `-vt` flag to make Epsilon run as a text program, or an appropriate `-display` flag to make Epsilon connect to a given X server. On platforms where Epsilon uses shared libraries, you can run the program `terminal-epsilon` instead of `epsilon`; it will run as a text program even where X11 shared libraries are not installed.

Epsilon also recognizes these standard X11 flags:

`-bw pixels` or `-borderwidth pixels` This flag sets the width of the window border in pixels. An `Epsilon.borderWidth` resource may be used instead.

`-display disp` This flag makes Epsilon use *disp* as the display instead of the one indicated by the `DISPLAY` environment variable. It follows the standard X11 syntax.

`-fn font` or `-font font` This flag specifies the font to use. The `Alt-x set-font` command can select a different font from within Epsilon. Epsilon will remember any font you select with `set-font` and use it in future sessions; this flag overrides any remembered font.

- geometry *geometry*** This flag sets the window size and position, using the standard X11 syntax for geometry flags. It uses the format `WIDTHxHEIGHT+XOFF+YOFF`. The `WIDTH` and `HEIGHT` values are in characters, separated by an “x”. The `XOFF` and `YOFF` values are in pixels, measured from the top left corner of the screen. You can use `-` instead of `+` as the offset separator to position relative to the right or bottom edge of the screen instead. You may omit trailing values (for instance, just specify width and height) and the offsets (following either a `+` or `-` separator) may be negative. Without this flag, Epsilon looks for an `Epsilon.geometry` resource, and if none, restores the size (and optionally, position: see `Epsilon.recordPosition` below) from the last time Epsilon exited.
- name *resname*** This flag tells Epsilon to look for X11 resources using a name other than “Epsilon”.
- title *title*** This flag sets the title Epsilon displays while starting. An `Epsilon.title` resource may be used instead.
- xrm *resourcestring*** This flag specifies a specific resource name and value, overriding any defaults.

Epsilon uses various X11 resources. You can set them from the command line with a flag like `-xrm Epsilon.cursorstyle:1` or put a line like `Epsilon.cursorstyle:1` in your X resources file, which is usually named `~/.Xresources` or `~/.Xdefaults`:

```
Epsilon.cursorstyle: 1
```

You’ll need to tell X to reread the file after making such a change, using a command like `xrdb -merge ~/.Xresources`.

As another option, you can put X resources in a file named `~/.epsilon/Xcustomresources` and Epsilon will load it each time it starts. These override any resources previously loaded on the server, and are in turn overridden by font selections you make within Epsilon (which Epsilon automatically saves), and by resources set on the command line.

Epsilon uses these X resources:

Epsilon.borderWidth This sets the width of the border around Epsilon’s window.

Epsilon.cursorstyle Under X11, Epsilon displays a block cursor whose shape does not change. Define a `cursorstyle` resource with value 1 and Epsilon will use a line-style cursor, sized to reflect overwrite mode or virtual space mode. Note: This cursor style does not display correctly on some X11 servers.

Epsilon.font This resource sets Epsilon’s font. It must be a fixed-width font. If you set a font from within Epsilon, it remembers your selection in a file `~/.epsilon/Xresources` and uses it in future sessions. Epsilon uses this resource if there’s no font setting in that file.

Epsilon.geometry This resource provides a geometry setting for Epsilon. See the `-geometry` flag above for the format.

Epsilon.recordPosition Define this resource with a value of 1 and Epsilon will remember and restore its window position on the screen, as well as its window size.

Epsilon.title This resource sets the title Epsilon displays while starting.

2.3 Installing Epsilon for macOS

Epsilon for macOS supports drag and drop installation. Simply open the downloaded disk image file and drag the Epsilon application inside to your Applications folder, or to the Applications alias. Epsilon supports macOS version 10.9 and later.

Epsilon for macOS runs as an X11 app when XQuartz is installed, and in console mode as a “Terminal” app when it isn’t. Installing XQuartz from <https://www.xquartz.org/> is highly recommended. It lets Epsilon display more colors, use font styles like bold and italic, recognize many more key combinations, and is also necessary if you want to edit files in folders like ~/Documents that receive privacy protection under macOS 10.15 Catalina and later. (See page 8 for more on macOS privacy contexts and Epsilon.)

Any time Epsilon documentation mentions the “Unix version” of Epsilon, this also includes the macOS version. Any time Epsilon documentation mentions Epsilon for X11, this includes Epsilon running under XQuartz (an X11 server). (For example, Epsilon for macOS running under XQuartz recognizes all the X11 flags described in the previous section, and all the X11 resource names documented there.)

2.3.1 Epsilon for macOS Keyboard Issues

When Epsilon runs under macOS, certain keyboard issues arise. This section explains how to resolve them.

- Under XQuartz, Epsilon uses the Command key as its Alt modifier key. XQuartz’s Preferences should be set so the “Enable key equivalents under X11” option is disabled (called “Enable Keyboard Shortcuts” in older XQuartz versions); otherwise XQuartz will reserve for itself many key combinations that use the Command key. Alternatively, you can substitute multi-key sequences like Escape f for the key combination Alt-f. See the alt-prefix command.
- macOS normally reserves the function keys F9 through F12 for its own use. Epsilon also uses these keys for various functions. You can set macOS to use different keys for these four functions system-wide, via System Preferences > Keyboard, but another option is to use alternative keys in Epsilon.

For the undo and redo commands on F9 and F10, the undo-changes and redo-changes commands on Ctrl-F9 and Ctrl-F10 make fine replacements. Or you can run undo and redo using their alternative key bindings Ctrl-x u and Ctrl-x r, respectively.

For the previous-buffer and next-buffer commands on F11 and F12, you can use their alternative key bindings, Ctrl-x < and Ctrl-x >, respectively.

- When Epsilon for macOS runs in console mode (typically, in an instance of the Terminal program) because XQuartz is not installed, it uses the TERM environment variable and the terminfo database of terminal characteristics. If you run Epsilon under a terminal program like Terminal and the TERM setting doesn’t match the terminal program’s actual behavior, some things won’t work right. The “xterm-color” setting comes closest to Terminal’s actual default behavior. Select this option from Terminal’s Preferences.

With the xterm-color setting, function keys F1-F4 may not work right; the commands on these keys almost all have alternative bindings you can use instead: For F1 (the help command), use the key labeled “Help” on Mac keyboards that have one, or type Alt-? or Ctrl-_. For F2 (the named-command command), use the Alt-x key combination instead. For F3 (the pull-word command), use the Ctrl-(Up) key. For F4 (the bind-to-key command), type Alt-x bind-to-key. Or you can change Terminal’s settings for these keys, or the terminfo database, so they match. But the best way to avoid these issues entirely is to install XQuartz so Epsilon can run as an X11 program, as above.

2.3.2 Epsilon for macOS Startup

When you run Epsilon for macOS, a helper program runs first and manages starting Epsilon. For example, it adds various flags to Epsilon’s command line that make it start in your home directory the very first time you run Epsilon, remember any directory you’ve set from session to session, and use a single instance.

Unlike most Unix-based systems, macOS applications run from an icon aren't descended from any shell process that loads user-defined environment variables. Since Epsilon uses some environment variables for configuration, on macOS, Epsilon's helper program loads environment variables from the file `~/epsilon/environment` if it exists. The format is a series of lines with `VAR=VALUE`, the same as the `env` command produces. Invalid lines are ignored. Any settings here override variables in the environment.

Epsilon includes a setup script, which will install Epsilon and its EEL compiler on your path, so you can run them from the command line more conveniently. Running it is optional. To run Epsilon's setup script from a shell prompt, type

```
sudo "/Applications/Epsilon 14.00.app/Contents/esetup"
```

assuming `/Applications` is where you installed Epsilon.

The setup script creates a link to Epsilon's startup helper, instead of running Epsilon's executable directly. If you want to run Epsilon using a shell alias, via an XQuartz Applications menu entry, or in some similar way, either have it run the startup helper program within Epsilon's app bundle (Epsilon 14.00.app/Contents/MacOS/epsilon-starter), or use the "open" command on Epsilon's app bundle, rather than directly running the Epsilon binary file within the app bundle, to let it switch privacy contexts if necessary. (See the next section.)

Normally Epsilon runs with a single instance when invoked via its icon, while each command-line invocation begins a new instance. You can change this by adding an entry to the file `~/epsilon/startup-config`. Add the line `singleinstance=always` to make Epsilon, when run from a command line, first try to bring any running instance to the front (via the `-add` flag), starting up normally if there's no previous instance running. Add the line `singleinstance=never` to make Epsilon's icon start a new Epsilon instance always. (This works for icons in a folder, not icons in the dock, which are always single-instance.)

Similarly, from an icon Epsilon normally restores the current directory from a previous session (via the `-w1` flag), while from the command line it inherits the invoking shell's current directory. You can change this by adding a line `useolddir=always` or `useolddir=never`, just as above. If you want the startup helper to pass a different flag in such cases, set `wflag=-w17`, for example. (The value 17 combines 1, setting the current directory from a previous session, and 16, overriding that by setting the current directory to that of the first file listed on the command line, if any.)

Some other settings are available to work around issues. To operate with the macOS privacy system explained in the next section, Epsilon's startup helper sometimes needs to know if X11's `DISPLAY` environment variable refers to a local display, where starting Epsilon in a different privacy context would work, or a remote ssh connection where it would not. At present, it checks to see if the value of `DISPLAY` contains "xquartz:", which indicates a local X11 server. You can change this string by setting `localdisplaydetect`, if that's not working right with your configuration. Epsilon also skips switching privacy contexts if invoked with `sudo`, but this check may be disabled with `ignoresudo=1`.

2.3.3 Epsilon for macOS and Privacy-protected Folders

When you run Epsilon for macOS, a helper program runs first and manages starting Epsilon. Starting in 10.15 Catalina, macOS uses a privacy system to keep programs from accessing privacy-protected folders like Downloads, Documents, or Desktop. Epsilon's startup helper tries to ensure it's running in a privacy context that permits macOS to ask for your permission before Epsilon can access these. (If not, macOS will either block Epsilon from accessing such folders, or apply the privacy context of some other program like Terminal.)

For this to work, XQuartz must be installed (so Epsilon runs as an X11 program), and you must either start Epsilon from its icon, or else (for command-line invocation), without using `sudo`. If you're connecting to your Mac over an ssh connection, that will also keep macOS from selecting Epsilon's privacy context.

You can change your privacy selections for Epsilon under System Preferences > Security & Privacy > Files and Folders, or let Epsilon access all files by manually adding it to System Preferences > Security & Privacy > Full Disk Access.

2.4 Installing Epsilon for DOS or Epsilon for OS/2

Older versions of Epsilon for DOS and OS/2 are also available for download.

To install Epsilon for DOS or OS/2, unzip the appropriate archive and run Epsilon's installation program by typing:

```
install
```

Follow the directions on the screen to install Epsilon. The installation program will ask before it modifies or replaces any system files. The DOS executable is named `epsdos.exe`, and the OS/2 executable is named `epsilon.exe`.

2.5 Tutorial

If you've never used Epsilon or EMACS before, you should run the tutorial to become acquainted with some of Epsilon's simpler commands.

The easiest way to run the tutorial is to start Epsilon and select Epsilon Tutorial from the Help menu. (If you're running a version of Epsilon without a menu bar, you can instead press the F2 key in Epsilon and type the command name `tutorial`. Or you can start Epsilon with the `-teach` flag.)

The tutorial will tell you everything else you need to know to use the tutorial, including how to exit the tutorial.

2.6 Invoking Epsilon

You can start Epsilon using the icon created by the installer. If you chose to put Epsilon on your PATH (for Windows, by selecting that option; for macOS, by running the `esetup` program), you can run Epsilon by simply typing `"epsilon"`.

Epsilon for Windows also includes a Windows Console version of Epsilon, which you can run by typing `"epsilonc"`. The corresponding terminal-based version on other platforms is called `"terminal-epsilon"`, or you can start Epsilon with its `-vt` flag. On Unix platforms, you can also run Epsilon using the command `"lugaru-epsilon"`, which is handy if your system includes a different and unrelated program also named `epsilon`.

The first time you run Epsilon, you will get a single window containing an empty document. You can give Epsilon the name of a file to edit on the command line. For example, if you type

```
epsilon sample.c
```

then Epsilon will start up and read in the file `sample.c`. If the file name contains spaces, surround the entire name with double-quote characters.

```
epsilon "a sample file.c"
```

When you name several files on the command line, Epsilon reads each one in, but puts only up to three in windows (so as not to clutter the screen with tiny windows). You can set this number by modifying the `max-initial-windows` variable.

If you specify files on the command line with wild cards, Epsilon will show you a list of the files that match the pattern in dired mode. See page 139 for more information on how dired works. File names that contain only extended wildcard characters like `;` `[` or `]`, and no standard wildcard characters like `*` or `?`, will be interpreted as file names, not file patterns. (If you set the variable `expand-wildcards` to 1, Epsilon will instead read in each file that matches the pattern, as if you had listed them explicitly. Epsilon for Unix does this too unless you quote the file pattern.)

Epsilon normally shows you the beginning of each file you name on the command line. If you want to start at a different line, put “+*number*” before the file’s name, where *number* indicates the line number to go to. You can follow the line number with a `:`column number too. For example, if you typed

```
epsilon +26 file.one +144:20 file.two
```

then you would get `file.one` with the cursor at the start of line 26, and `file.two` with the cursor at line 144, column 20. You can instead specify a character offset using the syntax “+*pnumber*” to go to character offset *number* in the buffer.

Windows users running the Cygwin environment may wish to configure Epsilon to accept Cygwin-style file names on the command line. See the `cygwin-filenames` variable for details.

By default, Epsilon will also read any files you were editing in your previous editing session, in addition to those you name on the command line. See page 144 for details.

If you’re running an evaluation version of Epsilon or a beta test version, you may receive a warning message at startup indicating that soon your copy of Epsilon will expire. You can disable or delay this warning message (though not the expiration itself). Create a file named `no-expiration-warning` in Epsilon’s main directory. Put in it the maximum number of days warning you want before expiration.

2.7 Configuration Variables: The Environment and The Registry

Epsilon for Unix uses several environment variables to set options and say where to look for files. Epsilon for Windows stores such settings in the System Registry, under the key `HKEY_CURRENT_USER\SOFTWARE\Lugaru\Epsilon`. Epsilon’s setup program will generally create all necessary registry keys automatically.

We use the term *configuration variable* to refer to any setting that appears as an environment variable under Unix, or a registry entry under Windows. There are a small number of settings that are stored in environment variables on all platforms; these are generally settings that are provided by the operating system. These include `COMSPEC`, `TMP` or `TEMP`, `EPSRUNS`, and `MIXEDCASEDRIVES`.

Under Windows, the installation program creates a registry entry similar to this:

```
HKEY_CURRENT_USER\SOFTWARE\Lugaru\Epsilon\EpsPath=~;c:\epsilon
```

Of course, the actual entry, whether it’s an environment variable setting or an entry in the system registry, would contain whatever directory Epsilon was actually installed in, not `c:\epsilon`.

If you have more than one version of Epsilon on your computer, you may want each to use a different set of options. You can override many of the configuration variables listed below by using a configuration variable whose name includes the specific version of Epsilon in use. For example, when Epsilon needs to locate its help file, it normally uses a configuration variable named `EPSPATH`. Epsilon version 6.01 would first check to see if a configuration variable named `EPSPATH601` existed. If so, it would use that variable. If not, it would then try `EPSPATH60`, then `EPSPATH6`, and finally `EPSPATH`. Epsilon does the same sort of thing with all the configuration variables it uses, with the exception of `DISPLAY`, `EPSRUNS`, `TEMP`, and `TMP`.

Epsilon uses a similar procedure to distinguish registry entries for the Windows Console mode version from registry entries for the Windows GUI version of Epsilon. For the console version, it checks registry names with an -NTCON suffix before the actual names; for the GUI version it checks for a -WIN suffix. So Epsilon 10.2 for Windows Console would seek an EPSPATH configuration variable using the names EPSPATH102-NTCON, EPSPATH102, EPSPATH10-NTCON, EPSPATH10, EPSPATH-NTCON, and finally EPSPATH, using the first one it finds.

For example, the Windows installation program for Epsilon doesn't actually add the EPSPATH entry shown above to the system registry. It really uses an entry like

```
HKEY_CURRENT_USER\SOFTWARE\Lugaru\Epsilon\EpsPath80=c:\epsilon
```

where EpsPath80 indicates that the entry should be used by version 8.0 of Epsilon, or version 8.01, or 8.02, but not by version 8.5. In this way, multiple versions of Epsilon can be installed at once, without overwriting each other's settings. This can be helpful when upgrading Epsilon from one version to the next.

Here we list all the configuration variables that Epsilon can use. Remember, under Windows, most of these names refer to entries in the registry, as described above. Under Unix, these are all environment variables.

CMDCONCURSHELLFLAGS If defined, Epsilon puts the contents of this variable before the command line when you use the start-process command with a numeric argument. It overrides CMDSHELLFLAGS. See page 150.

CMDSHELLFLAGS If defined, Epsilon puts the contents of this variable before the command line when it runs a subshell that should execute a single command and exit.

COMSPEC Epsilon for Windows needs a valid COMSPEC environment variable in order to run another program. Normally, the operating system automatically sets up this variable to give the file name of your command processor. If you change the variable manually, remember that the file must actually exist. Don't include command line options for your command processor in the COMSPEC variable. If a configuration variable called EPSCOMSPEC exists, Epsilon will use that instead of COMSPEC. (For Unix, see SHELL below.)

DISPLAY Epsilon for Unix tries to run as an X11 program if this environment variable is defined, using the X server display it specifies.

EEL The EEL compiler looks for a configuration variable named EEL before examining its command line, then "types in" the contents of that variable before the compiler's real command line. See page 397.

EPSCOMSPEC See COMSPEC above.

EPSCONCURCOMSPEC If defined, Epsilon for Windows runs the shell command processor named by this variable instead of the one named by the EPSCOMSPEC or COMSPEC variables, when it starts a concurrent process. See page 150.

EPSCONCURSHELL If defined, Epsilon for Unix runs the shell command processor named by this variable instead of the one named by the EPSSHELL or SHELL variables, when it starts a concurrent process. See page 150.

EPSCUSTDIR Epsilon uses the directory named here as its customization directory (see page 13) instead of the usual one (under \Users or \Documents and Settings, for Windows, or at ~/.epsilon, for Unix). The directory must already exist, or Epsilon will ignore this variable.

EPSILON Before examining the command line, Epsilon looks for a configuration variable named EPSILON and "types in" the value of that variable to the command line before the real command line. See page 13.

EPSMIXEDCASEDRIVES This variable can contain a list of drive letters. If the variable exists, Epsilon doesn't change the case of file names on the listed drives. See page 126 for details.

EPSPATH Epsilon uses this configuration variable to locate its files. See page 12.

EPSRUNS When Epsilon runs another program, it sets this environment variable to indicate to the other program that it's running within Epsilon. A setting of C indicates the subprocess is running within Epsilon's concurrent process. A setting of P indicates the subprocess is running via the filter-region command or similar. A setting of Y indicates Epsilon ran the process in some other way, such as via the shell command.

EPSSHELL See SHELL below.

ESESSION Epsilon uses this variable as the name of its session file. See page 144.

INTERCONCURSHELLFLAGS If defined, Epsilon uses the contents of this variable as the command line to the shell command processor it starts when you use the `start-process` command without a numeric argument. It overrides `INTERSHELLFLAGS`. See page 150.

INTERSHELLFLAGS If defined, Epsilon uses the contents of this variable as a subshell command line when it runs a subshell that should prompt for a series of commands to execute. See page 150.

MIXEDCASEDRIVES This variable can contain a list of drive letters. If the variable exists, Epsilon doesn't change the case of file names on the listed drives. See page 126 for details.

NOFOCUSCLICK If defined, when you click on an Epsilon window under MS-Windows, while another program has the focus, Epsilon will get the focus but will otherwise ignore the mouse click. By default, it treats mouse clicks the same whether or not they switch the focus to Epsilon, setting point to the character you clicked on.

PATH The operating system uses this variable to find executable programs such as `epsilon.exe`. Make sure this variable includes the directory containing Epsilon's executable files if you want to conveniently run Epsilon from the command line.

SHELL Epsilon for Unix needs a valid SHELL environment variable in order to run another program. If a configuration variable called EPSSHELL exists, Epsilon will use that instead of SHELL. (See COMSPEC above for the non-Unix equivalent.)

TEMP Epsilon puts any temporary files it creates in this directory, unless a TMP environment variable exists. See the description of the `-fs` flag on page 14.

TMP Epsilon puts any temporary files it creates in this directory. See the description of the `-fs` flag on page 14.

2.7.1 How Epsilon Finds its Files

Sometimes Epsilon needs to locate one of its files. For example, Epsilon needs to read an `.mnu` file like `gui.mnu` or `epsilon.mnu` to determine what commands go in its menu bar.

Epsilon searches for the file in each directory named by the EPSPATH configuration variable. This configuration variable should contain a list of directories, separated by semicolons (or for Unix, colons). Epsilon will then look for the file in each of these directories. Under Windows, a directory named `~` in an EPSPATH variable has a special meaning. It refers to the current user's customization directory. See the next section.

If there is no EPSPATH configuration variable, which is the usual case on non-Windows systems, Epsilon constructs a default one. It consists of the user's customization directory, then the parent of the directory containing Epsilon's executable. For Unix, the default EPSPATH also contains the directories `/opt/epsilonVER` and `/usr/local/epsilonVER` (where *VER* indicates the current version, such as 10.01). For Mac OS, it also contains the Misc and Resources subdirectories within its app bundle, just after your customization directory.

If the name of the directory with Epsilon's executable doesn't start with `bin`, or its parent doesn't start with `eps` (they do, in a normal installation), Epsilon uses the directory containing Epsilon's executable, not its parent, in the default EPSPATH. (MacOS doesn't use this rule unless you manually rearrange files within its app bundle.)

Some flags can change the above behavior. The `-w32` flag makes Epsilon look for files in the directory containing the Epsilon executable before trying the EPSPATH. The `-w8` flag keeps Epsilon from including the executable's directory or its parent in the default EPSPATH.

The EEL compiler also uses the EPSPATH configuration variable. See page 397.

2.7.2 The Customization Directory

Epsilon searches for some files in a user-specific customization directory. It also creates files like its initialization file `einit.ecm` there. (See page 164, and the `edit-customizations` command.)

To locate your customization directory, switch to Epsilon's `#messages#` buffer. Epsilon writes the name of its customization directory to this buffer when it starts up. Or run the `edit-customizations` command, which opens the `einit.ecm` file located in this directory.

Under Linux, FreeBSD, and macOS, the customization directory is `~/.epsilon`.

Under Windows, the customization directory is located in the `Lugaru\Epsilon` subdirectory within the current user's Application Data directory:

```
\Users\username\AppData\Roaming\Lugaru\Epsilon
```

or

```
%APPDATA%\Lugaru\Epsilon
```

On all platforms, you can use the syntax `%@ecustom%` in Epsilon to refer to the customization directory.

You can force Epsilon to use a different customization directory by defining a configuration variable named `EPSCUSTDIR`. See page 10 for more on configuration variables.

2.8 Epsilon Command Line Flags

When you start Epsilon, you may specify a sequence of command line flags (also known as command-line options, or switches) to alter Epsilon's behavior. Flags must go before any file names.

Each flag consists of a minus sign ("`-`"), a letter, and sometimes a parameter. You can use the special flag `--` to mark the end of the flags; anything that follows will be interpreted as a file name even if it starts with a `-` like a flag.

If a parameter is required, you can include a space before it or not. If a parameter is optional (`-b`, `-m`, `-p`) it must immediately follow the flag, with no space.

Before examining the command line, Epsilon looks for a configuration variable (see page 10) named EPSILON and “types in” the value of that variable to the command line before the real command line. Thus, if you define a Unix environment variable:

```
export EPSILON=-m250000 -smine
```

then Epsilon would behave as if you had typed

```
epsilon -m250000 -smine myfile
```

when you actually type

```
epsilon myfile
```

Here we list all of the flags, and what they do:

- +*number*** Epsilon normally shows you the beginning of each file you name on the command line. If you want to start at a different line, put “+*number*” before the file’s name, where *number* indicates the line number to go to. You can follow the line number with a colon and a column number if you wish.
- add** This flag tells Epsilon to locate an existing instance of Epsilon, pass it the rest of the command line, and exit. Epsilon ignores the flag if there’s no prior instance. If you want to configure another program to run Epsilon to edit a file, but use an existing instance of Epsilon if there is one, just include this flag in the Epsilon command line. See page 146 for details on Epsilon’s server support.
- b*filename*** Epsilon normally reads all its commands from a state file at startup. (See the -s flag below.) Alternately, you can have Epsilon start up from a file generated directly by the EEL compiler. These *bytecode files* end with a “.b” extension. This flag says to use the bytecode file with name *filename*, or “epsilon” if you leave out the *filename*. You may omit the extension in *filename*. You would rarely use this flag, except when building a new version of Epsilon from scratch. Compare the -l flag.
- d*variable!value*** You can use this flag to set the values of string and integer variables from the command line. The indicated variable must already exist at startup. You can also use the syntax -d*variable=value*, but beware: if you run Epsilon for Windows via a .BAT or .CMD file, the system will replace any =’s with spaces, and Epsilon will not correctly interpret the flag.
- dir *dirname*** Epsilon interprets any file names that follow on the command line relative to this directory.
- fd*filename*** This flag tells Epsilon where to look for the on-line documentation file. Normally, Epsilon looks for a file named edoc. This flag tells Epsilon to use *filename* for the documentation file. If you provide a relative name for *filename*, then Epsilon will search for it; see page 12. Use a file name, not a directory name, for *filename*.
- fs*dirname*** This switch tells Epsilon what directories to use for temporary files, such as Epsilon’s swap file, which it uses when you edit files too big for available memory, or the eshell file it creates in some environments to help capture the output of a process. *Dirnames* should indicate a list of one or more directories, separated by semicolons (colons under Unix). Epsilon will use the first directory named as long as there is space on its device; then it will switch to the second directory, and so forth. If it cannot find any available space, it will ask you for another directory name.

If you don’t use this switch, Epsilon will create any temporary files it needs in the directory named by the TMP environment variable. If TMP doesn’t exist, Epsilon tries TEMP, then picks a fallback location. Epsilon calls its swap file eswap, but it will use another name (like eswap0, eswap1, etc.) to avoid a conflict with another Epsilon using this file.
- geometry** When Epsilon for Unix runs as an X program, it recognizes this standard X11 flag, which specifies the size and position of Epsilon’s window. See page 6.

-k $anumber$ This switch changes certain keyboard and display functions, primarily to help diagnose problems. It's followed by a number, a bit pattern made by summing the bit values that follow.

For Windows, the value 1 tells Epsilon not to translate the Ctrl-2 key combination to Ctrl-@. (Ctrl-Shift-2 always produces Ctrl-@.) The value 16 makes text a little darker, and sometimes helps with display driver compatibility too.

A value of 128 tells Epsilon for Windows not to apply the Ctrl key to those ASCII characters that have no Control version in ASCII. For instance, the ASCII code includes characters Ctrl-a and Ctrl-\, but not Ctrl-9 or Ctrl-(. Epsilon for Windows will construct a non-ASCII key code for the latter pair unless you use this bit. (Under X11, Epsilon always does this.)

For Unix, bits in this flag can set which X11 modifier keys indicate an Alt key. By default, Epsilon chooses an appropriate key, but you can use 1 or 2 to force modifier key 1 or 2, respectively. The number is a bit pattern specifying which of the five possible X11 modifier keys will be used as an Alt key, using the values 1, 2, 4, 8, and 16. The value 32 tells Epsilon under X11 not to translate the Ctrl-2 key combination to NUL (as 1 for Windows does).

Both Windows and X11 GUI versions recognize the 64 bit, which tells Epsilon not to translate the Ctrl-6 combination into Ctrl-^, or Ctrl-(Minus) on the main keyboard into Ctrl-_-.

-ks $number$ This flag lets you adjust the emphasis Epsilon puts on speed during long operations versus responsiveness to the abort key. Higher numbers make Epsilon slightly faster overall, but when you press the abort key, Epsilon may not respond as quickly. Lower numbers make Epsilon respond more quickly to the abort key, but with a performance penalty. The default setting is -ks100.

-l $bytecode$ Giving this switch makes Epsilon load a bytecode file named *bytecode.b* after loading the state file. If you give more than one -l flag on the command line, the files load in the order they appear. Compare the -b flag.

-m $bytes$ This switch controls how much memory Epsilon uses for the text of buffers. Epsilon interprets a number less than 1000 as a number of kilobytes, otherwise, as bytes. You may explicitly specify kilobytes by ending *bytes* with 'k', or megabytes by ending *bytes* with 'm'. Specify -m0 to use as little memory as possible, and -m to put no limit on memory use.

If you read in more files than will fit in the specified amount of memory, or if despite a high limit, the operating system refuses Epsilon's requests for more memory, Epsilon will swap portions of the files to disk. By default, Epsilon puts no limits on its own memory usage.

-noinit This flag tells Epsilon not to read any *enit.ecm* customization file.

-nologo In some environments Epsilon prints a short copyright message when it starts. This flag makes it skip displaying that message.

-noserver This flag tells Epsilon for Windows or Unix that it should not register itself as a server so as to accept messages from other instances of Epsilon. By default, Epsilon will receive messages from future instances of Epsilon that are started with the -add flag, or (for Windows) sent via file associations or DDE. See page 146 for details. The flag -nodde is a synonym.

-p $filename$ This overrides the ESESSION configuration variable to control the name of the session file that Epsilon uses. When you specify a file name, Epsilon uses that for the session file, just as with ESESSION. Because the -p0 and -p1 flags enable and disable sessions (see the next item), the given *filename* must not begin with a digit.

-p $number$ This flag controls whether or not Epsilon restores your previous session when it starts up. By default, Epsilon will try to restore your previous window and buffer configuration. The -p flag with no number toggles whether Epsilon restores the session. Give the -p0 flag to disable session restoring

and saving, and the `-p1` flag to enable session restoring and saving. This flag understands the same values as the `preserve-session` variable; see its description for other options.

- quickup** Epsilon uses this flag to help perform certain updates. It searches for and loads a bytecode file named `quickup.b`. This flag is similar to the `-l` flag above, but the `-quickup` flag doesn't require any EEL functions to run. For that reason, it can replace and update any EEL function.
- rcommand** Giving this switch makes Epsilon try to run a command or keyboard macro named *command* at startup. If the command doesn't exist, nothing happens. If you specify more than one `-r` flag on the command line, they execute in the order they appear. Use the syntax `-rcmdname=param` or `-rcmdname!param` to run an EEL subroutine and pass it a value; the subroutine must be defined to accept a single parameter of `char *` type.
- sfilename** When Epsilon starts up, it looks for a *state file* named `epsilon-v14.sta`. The state file contains definitions for all of Epsilon's commands. You can create your own state file by using the `write-state` command. This switch says to use the state file with the name *filename*. Epsilon will add the appropriate extension if you omit it. Specify a file name for *filename*, **not** a directory name. Of course, the file name may include a directory or drive prefix. If you specify a relative file name, Epsilon will search for it. See page 12. See also the `-b` flag, described above.
- sendonly** The startup script in Epsilon for macOS uses this flag in combination with the `-add` flag. It makes Epsilon exit with an error code whenever no prior instance was found to receive the `-add` command line.
- server:servename** The command line flag `-server` may be used to alter the server name for an instance of Epsilon. An instance of Epsilon started with `-server:somename -add` will only pass its command line to a previous instance started with the same `-server:somename` flag. See page 146. The flag `-dde` is a synonym.
- teach** This flag tells Epsilon to load the on-line tutorial file at startup. See page 9.
- vcx** *x* indicates the number of columns you want displayed while in Epsilon. For example, use `"-vc132"` for 132 columns. See the `-vl` flag, described below. See the `-geometry` flag for the equivalent in Epsilon for Unix.
- vcolor** Epsilon normally tries to determine whether to use a monochrome color scheme or a full-color one based on the type of display in use and its mode. This flag forces Epsilon to use a full-color color scheme, regardless of the type of the display.
- vfnumber** This flag controls how Epsilon for Windows behaves at startup when the previous window position is partly or entirely offscreen. It's followed by a number, a bit pattern made by summing the bit values listed below.

By default, Epsilon first picks which monitor (on a multi-monitor system) overlaps the most with its recorded position, then shrinks its window so it's no larger than that monitor. Then, Epsilon moves its window until at least half of it is visible on that monitor, doing this separately in the horizontal and vertical directions.

The 1 bit makes Epsilon omit shrinking its window. (Using the `-vc` or `-vl` flags to set the number of columns or lines also overrides any window resizing in that direction.)

The 2 bit makes Epsilon omit moving its window to put more of it on the monitor. Using the `-vx` or `-vy` flags to explicitly position Epsilon's window also overrides Epsilon's normal window shifting in that direction.

Normally Epsilon only considers the working area of a monitor, excluding elements like its task bar. The 4 bit makes it perform its calculations using the entire monitor area.

The 8 bit changes Epsilon's movement rule, making it move Epsilon's window so it doesn't extend off the selected monitor at all.

- vlx** *x* indicates the number of screen lines you want to use while in Epsilon. Also See the **-vc** switch, described above. See **-geometry** for the equivalent in Epsilon for Unix.
- vmono** Epsilon normally tries to determine whether to use a monochrome color scheme or a full-color one based on the type of display in use and its mode. This flag forces Epsilon to use its monochrome color scheme, regardless of the type of the display.
- vt** (Unix only) This flag forces Epsilon to run as a curses-style terminal program, not an X11 program. By default Epsilon for Unix runs as an X program whenever an X display is specified (either through a DISPLAY environment variable or a **-display** flag), and a terminal program otherwise.
- vv** This flag instructs Epsilon to split the screen vertically, not horizontally, when more than one file is specified on the command line.
- vx and -vy** These flags let you specify the position of Epsilon's window in Epsilon for Windows. For example, **-vx20 -vy30** positions the upper left corner of Epsilon's window at pixel coordinates 20x30. See **-geometry** for the equivalent in Epsilon for Unix.
- wnumber** This flag controls several directory-related settings. Follow it with a number.
 - The **-w1** flag tells Epsilon to remember the current directory from session to session. Without this flag, Epsilon will remain in whatever current directory it was started from. Epsilon always records the current directory when it writes a session file; this flag only affects whether or not Epsilon uses this information when reading a session file.
 - The **-w2** and **-w4** flags have no effect in this version of Epsilon.
 - The **-w8** flag tells Epsilon not to look for its own files in the parent of the directory containing the Epsilon executable. See page 12.
 - The **-w16** flag tells Epsilon to set its current directory to the directory containing the first file named on its command line. If you edit files by dragging and dropping them onto a shortcut to Epsilon, you may wish to use this flag in the shortcut.
 - The **-w32** flag tells Epsilon to look for its own files in the directory containing the Epsilon executable before searching the EPSPATH. See page 12.
 - You can combine **-w** flags by adding their values together. For example, **-w9** makes Epsilon remember the current directory and exclude its executable's parent directory from the default EPSPATH. These **-w** flags are cumulative, so **-w1 -w8** works the same as **-w9**. Omitting the number discards all prior **-w** flags on the command line, so **-w9 -w -w32** acts like just **-w32**.
 - Program icons for Epsilon typically invoke it with **-w1** so that Epsilon remembers the current directory.
- wait** This flag tells Epsilon to locate an existing instance of Epsilon, pass it the rest of the command line, and wait for the user in that instance to invoke the **resume-client** command. (Epsilon ignores the flag if there's no prior instance.) If you want to configure another program to run Epsilon to edit a file, but use an existing instance of Epsilon, just include this flag in the Epsilon command line. See page 146 for details on Epsilon's server support.

2.9 File Inventory

Epsilon consists of the following files:

epsilon.exe The Epsilon for Windows executable program.

epsilonoc.exe The Epsilon executable program for Windows Console mode.

eel.exe Epsilon's compiler. You need this program if you wish to add new commands to Epsilon or modify existing ones.

eel_lib.dll Under Windows, Epsilon's compiler eel.exe requires this file. Epsilon itself also uses this file when you compile from within the editor.

icu* These files help provide Unicode support under Windows.

inherit.exe and inherit.pif Epsilon for Windows uses these files to execute another program and capture its output.

edoc.hlp This Windows help file provides help on Epsilon.

epshlp.dll Epsilon's WinHelp file communicates with a running copy of 32-bit Epsilon so it can display current key bindings or variable values and let you modify variables from the help file. It uses this file to do that.

sendeps.exe The Epsilon for Windows installer uses this file to help create desktop shortcuts and Send To menu entries. See page 148.

VisEpsil.dll Epsilon for Windows includes this Developer Studio extension that lets Developer Studio pass all file-opening requests to Epsilon.

mspellcmd.exe Epsilon's speller uses this helper program to get suggestions from the MicroSpell speller.

econfig.exe Epsilon for Windows runs this program when you use the configure-epsilon command.

mshelp2.vbs Epsilon for Windows uses this script to display MS Help 2 files.

fixpath2.exe Epsilon's Windows installer and configuration program use this program to add or remove directories from the system's PATH.

winpty32.exe, winpty64.exe and win-askpass.exe, run-ssh-agent.bat The secure shell (ssh) and secure file transfer (scp) features in Epsilon for Windows use these helper programs to interact with Cygwin's ssh program.

bscquery.exe Epsilon for Windows uses this program to help support .bsc source code browser files.

owitheps.dll This shell extension can be used to put an Open With Epsilon menu item on File Explorer's context menu, but by default Epsilon does this in a way that requires only setting registry entries, so this DLL is unused.

The installation program puts the following files in the main Epsilon directory, normally \Program Files\Eps14 under Windows and /opt/epsilon14.00 under Unix. (Epsilon for macOS keeps these files in various directories located within its app bundle, following Apple's requirements.)

epsilon-v14.sta This file contains all of Epsilon's commands. Epsilon needs this file in order to run. If you customize Epsilon, this file changes. The name includes Epsilon's major version.

original.sta This file contains a copy of the original version of epsilon-v14.sta at the time of installation.

edoc Epsilon's on-line documentation file. Without this file, Epsilon can't provide basic help on commands and variables.

info\epsilon.inf Epsilon's on-line manual, in Info format.

info\dir A default top-level Info directory, for non-Unix systems that may lack one. See Info mode for details.

lhelp* This directory contains files for the HTML version of Epsilon's documentation. The lhelp helper program reads them.

eteach Epsilon's tutorial. Epsilon needs this file to give the tutorial (see page 9). Otherwise, Epsilon does not need this file to run.

keychart.pdf A printable sheet listing most of Epsilon's default key assignments.

colclass.txt One-line descriptions of each of the different color classes in Epsilon. The set-color command reads this file.

brief.kbd The brief-keyboard command loads this file. It contains the bindings of all the keys used in Brief emulation, written in Epsilon's command file format.

epsilon.kbd The epsilon-keyboard command loads this file. It contains the standard Epsilon key bindings for all the keys that are different under Brief emulation, written in Epsilon's command file format.

epsilon.mnu, brief.mnu Non-GUI versions of Epsilon use one of these files to construct the menu bar.

gui.mnu, cua.mnu GUI versions of Epsilon use one of these files to construct the menu bar.

latex.env The tex-environment command in LaTeX mode (Alt-Shift-E) gets its list of environments from this file. You can add new environments by editing this file.

filter.txt A file defining the options for the filter control of Epsilon's Common File Open/Save dialogs under Windows.

readme.txt This file describes changes in recent versions of Epsilon. You can use the Alt-x release-notes command to read it.

epsout.e The import-customizations command uses this file.

uninstall.exe If you used the Windows-based installer, you can uninstall Epsilon by running this program.

install.log The Windows-based installer creates this file to indicate which files it installed.

***.h** The installation program copies a number of "include files" to the subdirectory "include" within Epsilon's main directory. These header files are used if you decide to compile an Epsilon extension or add-on written in its EEL extension language.

eel.h Epsilon's standard header file, for use with the EEL compiler.

codes.h Another standard header file, with numeric codes. The eel.h file includes this one automatically.

***.e** These files contain source code in EEL to all Epsilon's commands. The installation program copies them to the subdirectory "source" within Epsilon's main directory.

epsilon.e This file loads all the other files and sets up Epsilon.

samplemode.e This example file demonstrates various aspects of how to define a new mode.

makefile You can use this file, along with a "make" utility program, to help recompile the above Epsilon source files. It lists the source files and provides command lines to compile them.

The directory “changes” within Epsilon’s main directory contains files that document new features added in Epsilon 9 and earlier versions. See the online documentation for details on changes in more recent versions. Other files in this directory may be used to help incorporate old customizations, when updating from Epsilon 7 or earlier. See page 170 for information on updating to a new version of Epsilon.

Chapter 3

General Concepts



This chapter describes the framework within which the commands operate. The chapter entitled “Commands by Topic”, which starts on page 37, goes into detail about every Epsilon command.

If you have never used Epsilon before, you should run the tutorial now. This chapter discusses some general facilities and concepts used throughout Epsilon by many of the commands. You will find the discussion much clearer if you’ve used the tutorial, and have become accustomed to Epsilon’s general style.

To run the tutorial, start Epsilon and select Epsilon Tutorial from the Help menu. (You can also press the F2 key in Epsilon and type the command name `tutorial`, or start Epsilon with the `-teach` flag.)

3.1 Buffers

In Epsilon’s terminology, a *buffer* contains text that you can edit. You can think of a buffer as Epsilon’s copy of a file that you have open for editing. Actually, a buffer may contain a copy of a file, or it may contain a new “file” that you’ve created but have not yet saved to disk.

To edit a file, you read the file into a buffer, modify the text of the buffer, and write the buffer to the file. A buffer need not necessarily correspond to a file, however. Imagine you want to write a short program from scratch. You fire up Epsilon, type the text of the program into a buffer, then save the buffer to a file.

Epsilon does not place any limitation on the number of active buffers during an editing session. You can edit as many buffers at the same time as you want. This implies that you can edit as many files, or create as many files, or both, as you desire. Each document or program or file appears in its own buffer.

3.2 Windows

Epsilon displays your buffers to you in *windows*. You can have one window or many windows. You can change the number and size of windows at any time. You may size a window to occupy the entire display, or to occupy as little space as one character wide by one character high.

Each window can display any buffer. You decide what a window displays. You can always get rid of a window without worrying about losing the information the window displays: deleting a window does **not** delete the buffer it displays.

Each window displays some buffer, and several windows can each display the same buffer. This comes in handy if you want to look at different parts of a buffer at the same time, say the beginning and end of a large file.

A buffer exists whether or not it appears in some window. Suppose a window displays a buffer, and you decide to refer to another file. You can read that file into the current window without disturbing the old buffer. You peruse the new buffer, then return to the old buffer.

You may find this scheme quite convenient. You have flexibility to arrange your buffers however you like on the screen. You can make many windows on the screen to show any of your buffer(s), and delete windows as appropriate to facilitate your editing. You never have to worry about losing your buffers by deleting or changing your windows.

Epsilon has many commands to deal with buffers and windows, such as creating, deleting, and changing the size of windows, reading files into a buffer, writing buffers out to files, creating and deleting buffers, and much more. We describe these in detail in the chapter “Commands by Topic”, which starts on page 37.

3.3 Epsilon’s Screen Layout

To see what buffers and windows look like, refer to figure 3.1. This shows what the screen looks like with only one window. It shows what the screen looks like when you edit a file named `screen.1`.

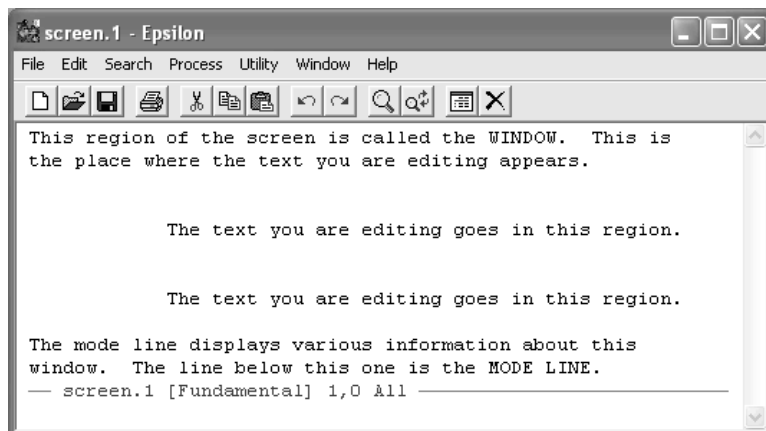


Figure 3.1: What Epsilon looks like with one window.

The top section of the screen displays some of the text of the window's buffer. Below that appears the *mode line*. The mode line begins with the name of the file shown in that buffer. If the buffer isn't associated with any file, Epsilon substitutes the buffer name, in parentheses.

Next comes the name of the current *major mode*, followed by any minor modes, all surrounded by square brackets. (See page 25.)

Then Epsilon shows the current column and line numbers (the first counting from zero, the second counting from 1), and the percentage of the buffer before the cursor. A star (*) at the end of the line means that you have changed the buffer since the last time you saved it to disk. (See the `mode-format` variable for information on customizing the contents of the mode line.) The text area and the mode line collectively constitute the window.

Below the mode line, on the last line of the screen, appears the *echo area*. Epsilon uses this area to prompt you for information or to display messages (in the figure it's empty). For example, the command to read a file into a buffer uses the echo area to ask you for the file name. Regardless of how many windows you have on the screen, the echo area always occupies the bottommost screen line.

When Epsilon displays a message in the echo area, it also records the message in the `#messages#` buffer (except for certain transient messages). See the `message-history-size` variable to set how Epsilon keeps the buffer from excessive size by dropping old messages.

Epsilon has an important concept called the editing point, or simply *point*. While editing a buffer, the editing point refers to the place that editing "happens", as indicated by the cursor. Point refers not to a character position, but rather to a character *boundary*, a place *between* characters. You can think of point as, roughly, the leftmost edge of the cursor. Defining the editing point as a position between characters rather than at a particular character avoids certain ambiguities inherent in the latter definition.

Consider, for example, the command that goes to the end of a word, `forward-word`. Since point always refers to a position between characters, point moves right after the last letter in the word. So the cursor itself would appear underneath the first character after the word. The command that moves to the beginning of the word, `backward-word`, positions point right before the first character in the word. In this case, the cursor itself would appear under the first character in the word.

When you want to specify a region, this definition for point avoids whether characters near each end belong to the region, since the ends do not represent characters themselves, but rather character boundaries.

Figure 3.2 shows Epsilon with 3 windows. The top window and bottom window each show the buffer "main". Notice that although these two windows display the same buffer, they show different parts of the




Figure 3.2: Epsilon with three windows.

buffer. The mode line of the top window says 0%, but the mode line of the bottom window says 58%. The middle window displays a different buffer, named “other”. If the cursor appears in the middle window and you type regular letters (the letters of your name, for example), they go into the buffer named “other” shown in that window. As you type the letters, the point (and so the cursor) stays to the right of the letters.

In general, the *current window* refers to the window with the cursor, or the window where the “editing happens”. The *current buffer* refers to the buffer displayed by the current window.

3.4 Different Keys for Different Uses: Modes

When you edit a C program, your editor should behave somewhat differently than when you write a letter, or edit a Lisp program, or edit some other kind of file.

For example, you might want the third function key to search forward for a comment in the current buffer. Naturally, what the editor should search for depends on the programming language in use. In fact, you might have PHP in the top window and C++ in the bottom window.

To get the same key (in our example, the third function key) to do the right thing in either window, Epsilon allows each buffer to have its own interpretation of the keyboard.

We call such an interpretation a *mode*. Epsilon comes with several useful modes built in, and you can add your own using the Epsilon Extension Language (otherwise known as EEL, pronounced like the aquatic animal).

Epsilon uses the mode facility to provide the *dired* command, which stands for “directory edit”. The *dired* command displays a directory listing in a buffer, and puts that buffer in dired mode. Whenever the current window displays that buffer, several special keys do things specific to dired mode. For example, the ‘e’ key displays the file listed on the current line of the directory listing, and the ‘n’ key moves down to the next line of the listing. See page 139 for a full description of dired mode.

Epsilon also provides C mode, which knows about several C indenting styles (see page 86) and is used for all C-like languages. Fundamental mode is a general-purpose editing mode used for scratch buffers and plain text files. And there are many other modes, some associated with specific commands (like hex mode, diff mode, or grep mode) and many more supporting individual programming languages or other file types. See the section starting on page 84.

Almost every mode has an associated command, named after the mode, that puts the current buffer in that mode. The c-mode and fundamental-mode commands put the current buffer into those modes, for instance.

Press F1 m to display help on the current buffer's major mode.

The mode name that appears in a mode line suggests the keyboard interpretation active for the buffer displayed by that window. When you start Epsilon with no particular file to edit, Epsilon uses Fundamental mode, so the word "Fundamental" appears in the mode line. Other words may appear after the mode name to signal changes, often changes particular to that buffer. We call these *minor modes*.

For example, the auto-fill-mode command sets up a minor mode that automatically types a `<Return>` for you when you type near the end of a line. (See page 77.) It displays "Fill" in the mode line, after the name of the major mode. A read-only buffer display "RO" to indicate that you won't be able to modify it. There is always exactly one major mode in effect for a buffer, but any number of minor modes may be active. Epsilon lists all active minor modes after the major mode's name.

Here are some common minor modes:

Fill indicates auto-filling is in effect for the current buffer. See page 77.

RO indicates the buffer is read-only. See page 119.

Pager is similar to RO, indicating the buffer is read-only and that `<Space>` and `<Backspace>` page forward and back, but this behavior isn't conditioned on the `readonly-pages` variable as read-only mode's is.

Def indicates Epsilon is defining a keyboard macro. See page 155.

Susp indicates defining or running a keyboard macro has been suspended. See page 156.

Narrow indicates only a section of the buffer is being displayed, and the rest has been hidden. See page 176.

Sp indicates Epsilon will highlight misspelled words in the current buffer. See page 81.

Along with any minor modes, Epsilon will sometimes also display the name of a type of file translation (one of DOS, Binary, Unix, or Mac). See page 122. It may also display the name of an encoding, such as UTF-8, OEM, or windows-1258. See page 136.

3.5 Keystrokes and Commands: Bindings

Epsilon lets you redefine the function of nearly all the keys on the keyboard. We call the connection between a key and the command that runs when you type it a *binding*.

For example, when you type the `<Down>` key, Epsilon runs the down-line command. The down-line command, as the name suggests, moves the point down by one line. So when you type the `<Down>` key, point moves down by one line.

You can change a key's binding using the bind-to-key command. The command asks for the name of a command, and for a key. Thereafter, typing that key causes the indicated command to run. Using bind-to-key, you could, for example, configure Epsilon so that typing `<Down>` would run the forward-sentence command instead of the down-line command.

This key-binding mechanism provides a great deal of flexibility. Epsilon uses it even to handle the alphabetic and number keys that appear in the buffer when you type them. Most of the alphabetic and number keys run the command normal-character, which simply inserts the character that invoked it into the buffer.

Out of the box, Epsilon comes with a particular set of key bindings that make it resemble the EMACS text editor that runs on many kinds of computers. Using the key-binding mechanism and the `bind-to-key` command, you could rearrange the keyboard to make it resemble another editor's keyboard layout. That is exactly what the `brief-keyboard` command does; it rearranges the keyboard commands to make Epsilon work like the Brief text editor. See page 159.

Epsilon provides over 400 commands that you can bind to keys, and you can write brand new commands to do almost anything you want, and assign them to whatever keys you choose. See page 157 for more information on the `bind-to-key` command.

Some commands have no default binding. You can invoke any command, bound or not, by giving its name. The command named `command`, normally bound to `Alt-x`, prompts for a command name and executes that command. For example, if you type

```
Alt-x down-line
```

followed by pressing the `<Enter>` key, the cursor moves down one line. Of course, you would find it easier in this example to simply type the `<Down>` key.

3.6 Repeating: Numeric Arguments

You can prefix a *numeric argument*, or simply an *argument*, to a command. This numeric argument generally functions as a repeat count for that command. You may enter a numeric argument in several ways. You may type `Ctrl-u` and then the number. You can also enter a numeric argument by holding down the `Alt` key and typing the number using the number keys across the *top* of the keyboard. Then you invoke a command, and that command generally repeats that number of times.

For example, suppose you type the four characters `Ctrl-u 26 Ctrl-n`. The `Ctrl-n` key runs the command named `down-line`, which moves point down one line. But given a numeric argument of 26, the command moves point down 26 lines instead of 1 line. If you give a numeric argument of -26 by typing a minus key while typing the 26, the `down-line` command would move point *up* 26 lines. You can get the same effect as `Ctrl-u 26 Ctrl-n` by holding down the `Alt` key and typing 26 on the main keyboard, then typing `Ctrl-n`. (Remember to release the `Alt` key first; otherwise you'd get `Alt-Ctrl-n`.)

You can give a numeric argument to any Epsilon command. Most commands will repeat, as our example did above. But some commands use the numeric argument in some other way, which can vary from command to command. Some commands ignore the numeric argument. We describe all the commands in the chapter titled "Commands by Topic", which starts on page 37.

3.7 Viewing Lists

Sometimes Epsilon needs to show you a list of information. For example, when it asks you for the name of a file to edit, you might request a list of possible files to edit (see the next section). In such cases, Epsilon will display the list of items in a pop-up window. While in a pop-up window, one line will stand out in a different color. If you press `<Enter>`, you select that item. To select another item, you can use normal Epsilon commands such as `<Up>` and `<Down>` to move to the next and previous items, or `<PageDown>` and `<PageUp>` to go to the next or previous windowful of items. You can even use Epsilon's searching commands to find the item you want. If you don't want any item on the list, you can simply type another response instead.

If you want to select one of the items and then edit it, press `Alt-e`. Epsilon will copy the highlighted line out of the list so can edit it.

3.8 Typing Less: Completion & Defaults

Whenever Epsilon asks you for some information (for instance, the name of a file you want to edit), you can use normal Epsilon commands to edit your response. For example, Control-A moves to the beginning of the response line. Most commands will work here, as long as the command itself doesn't need to prompt you for more information.

At many prompts, Epsilon will automatically type a default response for you, and highlight it. Editing the response will remove the highlight, while typing a new response will replace the default response. You can set the variable `insert-default-response` to zero if you don't want Epsilon to type in a response at prompts.

If you type a Control-R or Control-S, Epsilon will type in the default text. This is especially useful if you've told Epsilon not to automatically insert the default response, but it can also come in handy when you've mistakenly deleted or edited the default response, and you want to get it back. It's also convenient at prompts where Epsilon doesn't automatically type the default response, such as search prompts. Epsilon keeps separate defaults for the regular expression and non-regular expression replace commands, and for the regular expression and non-regular expression search commands. Epsilon will never overwrite what you actually type with a default, and indeed will only supply a default if you haven't yet specified any input for the response.

Another way to retrieve a previous response is to type Alt-e. While Ctrl-r and Ctrl-s provide a "suggested response" in many commands, Alt-e always types in exactly what you typed to that prompt last time. For example, at the prompt of the write-file command, Ctrl-s types in the name of the directory associated with the file shown in the current window, while Alt-e types in the last file name you typed at a write-file prompt. See page 30.

Alt-g provides yet another suggested response; it's often the name of the "current thing" for this prompt; in a search-and-replace command, for instance, Alt-g when typing the replacement text inserts the search text. In the write-file example, Alt-g inserts the current name of the file.

Sometimes Epsilon shows you the default in square brackets []. This means that if you just press (Enter) without entering anything, Epsilon will use the value between the square brackets. Often you can use the Ctrl-s or Alt-e keys to pull in that value, perhaps so that you can use regular Epsilon commands to edit the response string.

Epsilon can also retrieve text from the buffer at any prompt. Press the Alt-(Down) key or Alt-Ctrl-n to grab the next word from the buffer and insert it in your response. Press the key again to retrieve successive words. This is handy if there's a file name in the buffer that you now want to edit, for example. The keys Alt-(PageDown) or Alt-Ctrl-v behave similarly, but retrieve from the current position to the end of the line.

You can also use pull completion to retrieve text at a prompt that isn't at the current position, but elsewhere in the buffer. Begin typing the word you want to retrieve; then press Ctrl-(Up) (or Ctrl-(Down)) to grab the previous (or next) word in the buffer that starts with what you've typed. F3 is the same as Ctrl-(Up). See page 101 for details.

Whenever Epsilon asks for the name of something (like the name of a command, file, buffer, or tag), you can save keystrokes by performing *completion* on what you type. For example, suppose you type Alt-x to invoke a command by name, then type the letter 'v'. Only one command begins with the letter 'v', the visit-file command. Epsilon determines that you mean the visit-file command by examining its list of commands, and fills in the rest of the name. We call this process *completion*.

To use completion, type a (Space) and Epsilon will fill in as much of the name as possible. The letters Epsilon adds will appear as if you had typed them yourself. You can enter them by typing (Enter), edit them with normal editing commands, or add more letters. If Epsilon cannot add any letters when you ask for completion, it will pop up a list of items that match what you've typed so far. To disable automatic pop-ups on completion, set the `completion-pops-up` variable to zero.

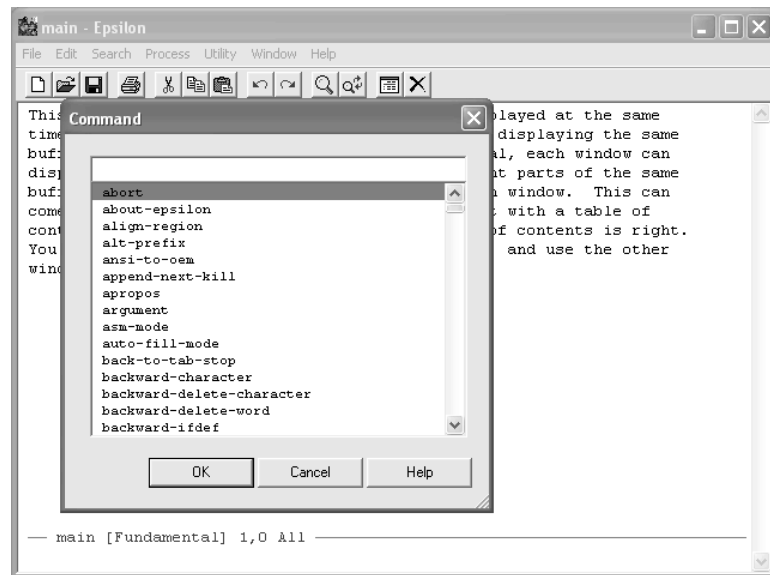


Figure 3.3: Typing ‘?’ shows all of Epsilon’s commands.

For example, four commands begin with the letters “go”, goto-beginning, goto-end, goto-line, and goto-tag. If you type “go”, and then press `<Space>`, Epsilon fills in “goto-” and waits for you to type more. Type ‘b’ and another `<Space>`, to see “goto-beginning”. Epsilon moves the cursor one space to the right of the last letter, to indicate a match. Press `<Enter>` to execute the goto-beginning command.

The `<Esc>` key works just like the `<Space>` key, except that if a single match results from the completion, Epsilon takes that as your response. This saves you a keystroke, but you don’t have the opportunity to check the name before continuing. The `<Tab>` key does the same thing. (These keys work the same in most Windows dialogs Epsilon displays, instead of performing their usual Windows functions of canceling or moving around in the dialog, respectively, but see the want-gui-prompts variable.)

Typing a question mark during completion causes Epsilon to display a list of choices in a pop-up window. Recall that completion works with buffer and file names, as well as with command names. For example, you can get a quick directory listing by giving any file command and typing a question mark when asked for the file name. Press the `Ctrl-g` key to abort the command, when you’ve read the listing. (See the dired command on page 139 for a more general facility.)

Figure 3.3 shows you what Epsilon looks like when you type `Alt-x` (the named-command command), and then press ‘?’ to see a list of the possible commands. Epsilon shows you all its commands in a pop-up window. Epsilon provides many more commands than could fit in the window, so Epsilon shows you the first window-full. At this point, you could press `<Space>` or `<PgDn>` to see the next window-full of commands, or use searching or other Epsilon commands to go to the item you desire. If you want the highlighted item, simply press `<Enter>` to accept it. If you type `Alt-e`, Epsilon types in the current item and allows you to edit it. Type any normal character to leave the pop-up window and begin entering a response by hand.

Figure 3.4 shows what the screen looks like if you type ‘w’ after the `Alt-x`, then type ‘?’ to see the list of possible completions. Epsilon lists the commands that start with ‘w’.

You can set variables to alter Epsilon’s behavior. The menu-width variable contains the width of the pop-up window of matches that Epsilon creates when you press ‘?’. (Unix only. In Windows, drag the dialog’s border to change its size.) The search-in-menu variable controls what Epsilon does when you press ‘?’ and then continue typing a response. If it has a value of zero, as it does by default, Epsilon moves from the pop-up window back to the response area, and editing keys like `<Left>` navigate in the response. If




Figure 3.4: Typing “w?” shows all commands that start with ‘w’.

`search-in-menu` has a nonzero value, Epsilon moves in the pop-up menu of names to the first name that matches what you’ve typed, and stays in the pop-up window. (If it can’t find a match, Epsilon moves back to the prompt as before.)

During file name completion, Epsilon can ignore files with certain extensions. The `ignore-file-extensions` variable contains a list of extensions to ignore. By default, this variable has the value `‘|.obj|.exe|.o|.b|’`, which makes file completion ignore files that end with `.obj`, `.exe`, `.o`, and `.b`. Each extension must appear between ‘|’ characters. You can augment this list using the `set-variable` command, described on page 159.

Similarly, the `only-file-extensions` variable makes completion look only for files with certain extensions. It uses the same format as `ignore-file-extensions`, a list of extensions surrounded by | characters. If the variable holds a null pointer, Epsilon uses `ignore-file-extensions` as above. Completion also restricts its matches using the `ignore-file-basename` and `ignore-file-pattern` variables, which use patterns to match the names of files to be excluded. When the pattern the user types doesn’t match any files due to such exclusions, Epsilon temporarily removes exclusions and lists matching files again.

When typing a file name, you can use environment variables to specify the directory or directories, using the syntax `%TEMP%` under Windows, and `$TEMP` or `${TEMP}` on macOS, Linux and FreeBSD. You can also use the values of Windows shell folders, and define your own shorthand names. See the `file-interpret-env-vars` variable for all the details.

3.9 Command History

Epsilon maintains a list of your previous responses to all prompts. To select a prompt from the list, press the `Alt-(Up)` key or `Alt-Ctrl-p`. Then use the arrow keys or the mouse to choose a previous response, and press `(Enter)`. If you want to edit the response first, press `Alt-e`.

For example, when you use the `grep` command to search in files for a pattern, you can press `Alt-(Up)` to see a list of file patterns you’ve used before. If the pattern `\windows\system*.inf` appeared on the list,

you could position the cursor on it and then press Alt-e. Epsilon would copy the pattern out of the list so you can edit it, perhaps replacing *.inf with *.ini. Both patterns would then appear in the history list next time. Or you could just press (Enter) in the list of previous responses to use the same pattern.

You can also use Alt-e at any prompt to retrieve the last response without showing a list of responses first. For example, Ctrl-x Ctrl-f Alt-e will insert the full name of the last file you edited with the find-file command.

Except in certain searching commands, you can press (Up) or Ctrl-p instead of Alt-(Up) key or Alt-Ctrl-p. These normally behave the same, but you can set the `recall-prior-response-options` variable to make the non-Alt versions of the keys select older command history responses without displaying a list of all of them.

3.10 Mouse Support

Epsilon supports a mouse under Windows and under X11 in Unix. You can use the left button to position point, or drag to select text. Double-clicking selects full words. (When a pop-up list of choices appears on the screen, double-clicking on a choice selects it.) Use shift-clicking to extend or contract the current selection by repositioning the end of the selection. Holding down the Alt key while selecting produces a rectangle selection.

Once you've selected a highlighted region, you can drag it to another part of the buffer. Move the mouse inside the highlighted region, hold down a mouse button and move the mouse to another part of the buffer while holding down the button. The mouse cursor changes to indicate that you're dragging text. Release the mouse button and the text will move to the new location. To make a copy of the text instead of moving it, hold down the Control key while dropping the text.

Dragging text with the mouse also copies the text to a kill buffer, just as if you had used the corresponding keyboard commands to kill the text and yank it somewhere else. When you drag a highlighted rectangular region of text, Epsilon's behavior depends upon the whether or not the buffer is in overwrite mode. In overwrite mode, Epsilon removes the text from its original location, replacing it with spaces. Then it puts the text in its new location, overwriting whatever text might be there before. In insert mode, Epsilon removes the text from its original location and shifts text to its right leftwards to fill the space it occupied. Then it shifts text to the right in the new location, making room for the text.

You can use the left button to resize windows by dragging window corners or borders. For pop-up windows only, dragging the title bar moves the window.

A pop-up window usually has a scroll bar on its right border. Drag the box or diamond up and down to scroll the window. Click on the arrows at the top or bottom to scroll by one line. Click elsewhere in the scroll bar to scroll by a page. In some environments, ordinary tiled windows have a scroll bar that pops up when you move the mouse over the window's right-hand border, or (for windows that extend to the right edge of the screen), when you move the mouse past the right edge. The `toggle-scroll-bar` command toggles whether tiled windows have pop-up scroll bars or permanent scroll bars.

Under X11, you can adjust the speed at which Epsilon scrolls due to mouse movements by setting the `scroll-rate` variable. It contains the number of lines to scroll per second. The `scroll-init-delay` variable contains the delay in hundredths of a second from the time the mouse button goes down and Epsilon scrolls the first time, to the time Epsilon begins scrolling repeatedly.

The right mouse button displays a context menu (which you can modify by editing an .mnu file, see page 33).

You can click (or hold) the middle mouse button and drag the mouse to pan or auto-scroll—the speed and direction of scrolling varies as you move the mouse. This works on wheeled mice or on any mouse with

three or more buttons. When you click the middle mouse button while holding down the Shift key, Epsilon pastes text instead. See the `mouse-center-yanks` variable to change its behavior.

Some mice have additional buttons. If the environment supports it, Epsilon recognizes buttons 4 and 5 and makes them scroll up and down. You can bind these buttons to different commands or to nothing using the `bind-to-key` and `unbind-key` commands.

Epsilon for Windows or Unix (under X11) also recognizes wheel rolling on wheeled mice, and scrolls the current window when you roll the wheel. See the `wheel-click-lines` variable for more details.

Under X11, some programs automatically make any text you select using the mouse available to be pasted in other programs. See the variable `mouse-selection-copies` to turn on this behavior for Epsilon.

3.11 The Menu Bar

The Windows GUI version of Epsilon provides a customizable menu bar and tool bar. To modify the menu bar, edit the file `gui.mnu`. See the next section for details. You can turn it off by adding `(set-gui-menu 0)` to your `einit.ecm` file (see page 164). To modify the tool bar, you can redefine the EEL command `standard-toolbar` in the file `menu.e`.

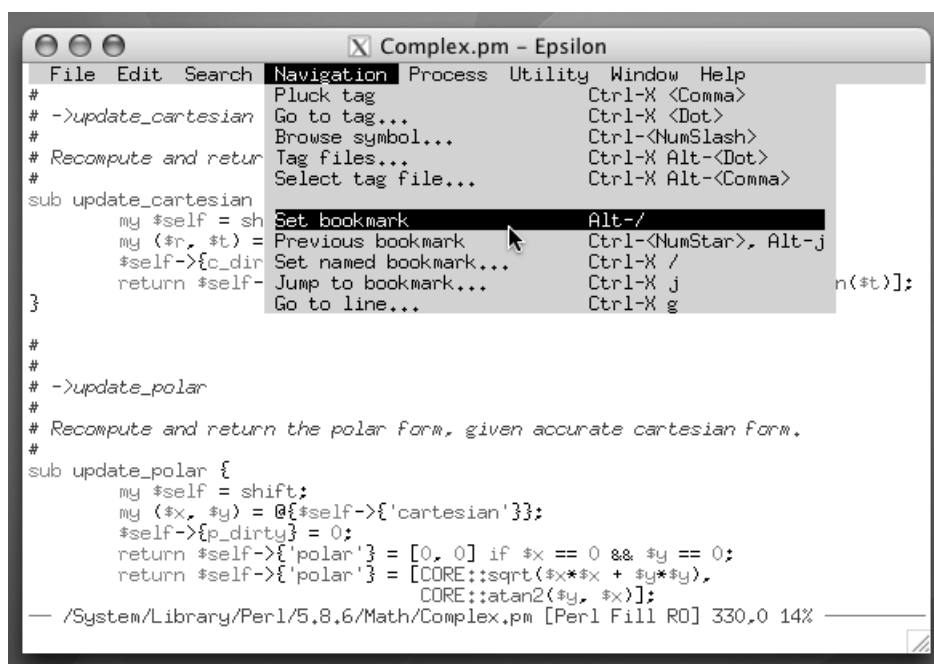


Figure 3.5: Epsilon's text-based menu bar.

Other versions of Epsilon provide a text-based menu bar, which is hidden by default. Most of the customization variables described below only apply to the text-based menu bar.

You can have Epsilon display a menu bar all the time with the `toggle-menu-bar` command, or press `Alt-F2` (the `show-menu` command) to display it at any time, and hide it again after you select a command. When you use the menu bar to invoke a command that needs additional input, Epsilon automatically brings up a list of options (as if you typed `?`) so that you can select one without using the keyboard.

You can change the contents of the menu bar by editing the appropriate `.mnu` file. See the next section.

If you hold down the Shift or Ctrl keys while selecting a menu command, Epsilon will run the command with a numeric argument of 1. This is handy for commands that behave differently when given a numeric argument. When you select an item on the text-based menu bar, Epsilon flashes the selected item. The `menu-bar-flashes` variable holds the number of flashes (default two).

By default, Epsilon displays key bindings for menu items. Set the variable `menu-bindings` to zero to disable this feature. (Epsilon for Windows ignores this variable and always displays such bindings.) Epsilon computes bindings dynamically the first time it displays a particular menu column. (For several commands with multiple bindings, the menu file selects a particular binding to display.) The `rebuild-menu` command makes Epsilon reconstruct its menus: use this command after setting `menu-bindings` or editing and saving a menu file.

By default, when you click on the text-based menu bar but release the mouse without selecting a command, Epsilon leaves the menu displayed until you click again. Set the `menu-stays-after-click` variable to zero if you want Epsilon to remove the menu when this happens.

3.11.1 Customizing Epsilon's Menu

You can change the contents of Epsilon's menu bar by editing a menu file, which uses an `.mnu` extension. Epsilon stores the name of its menu file in the variable `menu-file`, except for Epsilon for Windows which uses the variable `gui-menu-file` instead. Set the appropriate variable to make Epsilon use a different menu file.

Emulations for Brief and CUA set these variables to make Epsilon use an alternative menu suitable for those emulations, but by default, Epsilon for Windows uses a file named `gui.mnu`, while all other versions use the file `epsilon.mnu`. If you put a customized version of an `.mnu` file in your customization directory (see page 13), Epsilon will use it instead of the factory version.

The first line of a menu file holds the main menu bar. Each menu entry must have spaces on both sides. Each of the submenus that follow begins with a line that has the submenu title from the main menu (again, with spaces on both sides), then the width in characters of the submenu to create, not including command bindings. A submenu named `_popup` defines the right-click context menu. The individual menu entries follow, each line containing the menu item name, one or more tab characters, and the definition (normally the name of an Epsilon command to execute). A line starting with a tab puts a blank line in the menu. An actual blank line ends the submenu. A line starting with `#` is a comment.

If an entry contains a binding (meaning that text before the first tab character extends past the column width indicated for that submenu), Epsilon uses the binding text as-is. Otherwise, Epsilon adds bindings when it first displays the submenu.

```
Open in Notepad          %notepad "%f"
```

A menu item usually specifies the name of an Epsilon command to run (actually any EEL function that takes no parameters will work), but you can instead put `%` followed by the command line for an external program Epsilon should run. The command line is interpreted as a file name template, which means you can pass the the name of the current file name, or parts of the name, using `%` sequences, like `%f` for the full file name. (See page 122.)

```
Search the web           !"c:\path\to\chrome.exe" https://google.com
                        !"c:\path\to\chrome.exe" https://google.com/search?q=!
# Note: Above is actually one long line.
```

Or instead of the name of an Epsilon command, you can put the full path to an external program to run, surrounded by `!` characters. Optionally, following the second `!` character, you can put a second command

line with a third ! within it. If there's a highlighted region when the menu item is run, Epsilon will use this second command line and substitute the text of the highlighted region for the third ! character.

Or instead of an Epsilon command name, you can put the name of a Windows help file with a \$ character before it. Epsilon for Windows will then display that help file.

Chapter 4

Commands by Topic



This chapter lists all the Epsilon commands, grouped by topic. Each section ends with a summary of the keys, and the names you would use to invoke the commands by name, or to rebind them to other keys.

4.1 Getting Help

You can get help on Epsilon by typing F1, the *help* key. The help key will provide help at any time. If you type it during another command, help simply pops up a description of that command. Otherwise, the help command asks you to type an additional key to indicate what sort of help you want. Many of these options are also available directly from Epsilon's Help menu item, in versions with a menu bar.

The help command actually uses various commands which you can invoke individually. Here are the keys you can use at the help prompt.

Pressing **A** invokes the `apropos` command, which asks for a string, looks through the one-line descriptions of all the commands and variables, then pops up a list of commands or variables (and their descriptions) that contain the string, along with their key bindings. Highlighted words are links to the full documentation.

(The Info, HTML-based, and WinHelp formats of Epsilon's full manual each include their own search function. These will perform full-text searches throughout Epsilon's manual, often finding many more matches than `apropos` finds by searching one-line descriptions.)

Help's **K** option invokes the `describe-key` command. It prompts for a key and provides full documentation on what that key does.

The **C** option invokes the command `describe-command`, which provides full documentation on the command whose name you specify, and also tells which keys invoke that command.

The **B** option invokes the command `show-bindings`, which asks for a command name and gives you the keys that run that command.

The **I** option invokes the command `info`, which starts Info mode. Info mode lets you read the entire Epsilon manual, as well as any other documentation you may have in Info format. See page 39.

The **F** option is a shortcut into Epsilon's manual in Info mode. It prompts for some text, then looks up that text in the index of Epsilon's online manual. Just press `<Enter>` to go to the top of the manual. This option invokes the command `epsilon-info-look-up`; the command `epsilon-manual-info` goes to the top of Epsilon's documentation without prompting.

The **Ctrl-c** option prompts for the name of an Epsilon command, then displays an Info page from Epsilon's online manual that describes the command.

The **Ctrl-k** option prompts for a key, then displays an Info page from Epsilon's online manual that describes the command it runs.

The **Ctrl-v** option prompts for an Epsilon variable's name, then displays an Info page from Epsilon's online manual that describes that variable.

The **H** option displays Epsilon's manual in HTML format, by running a web browser. It prompts for a topic, which can be a command or variable name, or any other text. (The browser will try to find an exact match for what you type; if not, it will search for web pages containing that word.) When you're looking at Epsilon's manual in Info mode, using one of the previous commands, this command will default to showing the same topic in a browser.

The **W** option, in Epsilon for Windows versions prior to Vista, displays Epsilon's WinHelp help file. Like the Info-format manual, it contains the complete text of the Epsilon manual. (Windows Vista and later no longer support WinHelp natively, so such users should select HTML or Info help formats instead. It's possible to install WinHelp onto later Windows systems, then enable its macros with a registry entry, if you prefer this style of help. See Microsoft's web site.)

The **Q** option invokes the command `what-is`, which asks for a key and tells you what command would run if you typed that key.

The **R** option invokes the `describe-variable` command, which asks for a variable name and displays the help on that variable.

The **L** option invokes the `show-last-keys` command, which pops up a window that displays the last 60 keystrokes you typed.

The **M** option displays help on the major mode of the current buffer. For example, when you're editing a C file, this command displays help on C mode.

The **T** option shows context-sensitive help on the word at point. The help source varies based on the buffer's mode. See page 102.

The **V** command displays Epsilon's version number and similar information.

The **?** option displays information on the help command itself, including its options, just as typing the help key again would.

The **B** and **Q** options tell you about bindings without showing you the associated documentation on the command. In contrast to the first three options, these two display their information in the echo area, instead of popping up a window.

The `wall-chart` command creates a table showing the commands invoked by all the keys. It builds a chart in a buffer named "wall". The wall chart includes any changes you may have made to the normal key bindings. You can print it and attach it to any convenient wall using the `print-buffer` command.

Epsilon's help system keeps track of any changes that you make to Epsilon. For example, if you completely remap the keyboard, Epsilon's help system will know about it and still give you correct key binding information. And Epsilon's help system will also keep track of any commands or keyboard macros that you write and add to Epsilon.

The `release-notes` command reads and displays the release notes for this version of Epsilon.

Some of Epsilon's help commands use the on-line documentation file, `edoc`. This file contains descriptions for each of Epsilon's commands and variables. See the description of the `-fd` flag on page 14.

While some help commands provide help using a specific format like WinHelp or HTML help, others change their format based on the current platform. For instance, pressing the help key at a prompt shows help using WinHelp on earlier Windows systems, using HTML help on Vista and later Windows version (which don't include WinHelp) and on Unix under X11, and using popup help windows on Unix systems in console mode. You can select a different preferred help format by setting the variables `epsilon-help-format-win-gui`, `epsilon-help-format-win-console`, and `epsilon-help-format-unix-gui`.

Summary:	F1, Alt-?, Ctrl-_	help
	F1 A	apropos
	F1 K	describe-key
	F1 C	describe-command
	F1 R	describe-variable
	F1 L	show-last-keys
	F1 Q, F6	what-is
	F1 B, F5	show-bindings
	F1 Ctrl-c	info-goto-epsilon-command
	F1 Ctrl-k	info-goto-epsilon-key
	F1 Ctrl-v	info-goto-epsilon-variable
	F1 V	about-epsilon

F1 F

epsilon-info-look-up
 wall-chart
 release-notes
 epsilon-manual
 epsilon-manual-info

4.1.1 Info Mode

Epsilon's Info mode lets you read documentation in Info format. You can press F1 i to start Info mode. One example of documentation available in Info format is Epsilon's manual.

An Info document is divided into nodes. Each node describes a specific topic. Nodes are normally linked together into a tree structure.

Every node has a name, which appears on the very first line of the node. The first line might look like this:

```
File: cp, Node: Files, Next: More Options, Prev: Flags, Up: Top
```

That line also indicates that the node named "More Options" comes next after this "Files" node. And it says which node comes before it, and which node is its parent. (Some nodes don't have a "Next" or a "Prev" or an "Up" node.) In Info mode, the keys N, P, and U move to the current node's Next node, its Prev node, or its Up node (its parent node).

You can scroll through a node with the usual Epsilon commands, but Info mode also lets you use `<Space>` to page forward and `<Backspace>` to page back. When you're at the end of a node, the `<Space>` key goes on to the next one, walking the tree structure so you can read through an entire Info file. The `<Backspace>` key does the reverse; it goes to the previous node when you press it and you're already at the top of a node. (The keys `]` and `[` move ahead and back similarly, but don't page; use them when you don't want to see any more of the current node.)

Some nodes have menus. They look like this:

```
* Menu:

* Buffers::
* Flags::
* Switches: Flags.
```

Press the M key to select an item from a menu, then type the name of the item (the part before the `:` character). You can press `<Space>` to complete the name, or type just part of the name. The first two menu items let you type Buffers or Flags and go to a node with that same name; the last item lets you type Switches but Epsilon will go to a node named Flags.

You can also press a digit like 1, 2, 3 to go to the corresponding node in the current node's menu. Press 0 to go to the last node, whatever its number. So in the menu above, either 3 or 0 would go to the Flags node. Typically when you select a node from a menu, that node's Up will lead back to the node with the menu.

A node can also have cross-references. A cross-reference looks like this: `*Note: Command History::`. Use the F key to follow a cross reference; it completes like M does.

Instead of typing M or F followed by a node name, you can use `<Tab>` and `<Backtab>` to move around in a node to the next or previous menu item or cross-reference, then press `<Enter>` to follow it. Or you can double-click with the mouse to follow one.

Epsilon keeps a history of the Info nodes you've visited, so you can retrace your steps. Press L to go to the last Info node you were at before this one. Press L repeatedly to revisit earlier nodes. When you're done looking at Info documentation, press Q to exit Info mode.

Info documentation is tree-structured. Normally each separate program has its own file of documentation, and the nodes within form a tree. Each Info file normally has a node named "top" that's the top node in its tree. Then all the trees are linked together in a directory file named "dir", which contains a menu listing all the available files. The T key goes to the top node in the current file. The D key goes to the top node in the directory file. The `wrap-info-mode` variable controls how long lines display.

When a node name reference contains a word in parentheses, like (epsilon)Language Modes, it indicates the node is in a file whose name is inside the parentheses. (Otherwise the node must be in the current file.) If you omit the node name and just say (epsilon), the Top node is implied.

When a complete path to an Info file isn't specified (as is usually the case), Epsilon looks along an Info path. First it looks in each directory of the colon-separated list in the variable `info-path-unix` (or, in non-Unix versions of Epsilon, the semicolon-separated list in `info-path-non-unix`). These paths may use `%x` to indicate the directory containing Epsilon's executable. If the Info file still isn't found, Epsilon tries directories listed in any `INFOPATH` environment variable.

Press S to search in an Info file. You can use the same keys as in other Epsilon search commands to perform a regular expression search, word search, or control case folding. This command will jump from node to node if necessary to find the next match. If you use normal searching keys like Ctrl-s or Ctrl-r, they will report a failing search if there are no more matches in the current node. Press Ctrl-s or Ctrl-r again to have Epsilon continue the search into other nodes.

Press I to use an Info file's index. I <Enter> simply moves to the first index node in a file. Or you can type some text, and Epsilon will display each of the nodes in the file that have an index entry containing that text. Use <Comma> to advance to the next such entry.

There are a few more Info commands. B goes to the beginning of the current node, like Alt-<. > goes to the last node of the file, viewed as a hierarchy. G prompts for the name of a node, then goes there. (You can use it to reach files that might not be linked into the Info hierarchy.) H displays this documentation. And ? displays a short list of Info commands.

You can navigate to Epsilon's manual using Info commands, as explained above, but Epsilon also provides some shortcut commands. Press F1 Ctrl-c to look up an Epsilon command's full documentation by command name. Press F1 Ctrl-k, then press any key and Epsilon will show the documentation for whatever command it runs. Press F1 Ctrl-v to look up a variable. Press F1 f <Enter> to go to the top of Epsilon's documentation tree, or type a topic name before the <Enter> and Epsilon will look up that word in the index to Epsilon's online documentation.

If you write your own Info file, Epsilon provides some commands that help. The `info-validate` command checks an Info file for errors (such as using a nonexistent node name). The `info-tagify` command builds or updates an Info file's tag table. (Info readers like Epsilon can find nodes more quickly when a file's tag table is up to date, so run this after you modify an Info file.)

Summary:	Info mode only: N	info-next
	Info mode only: P	info-previous
	Info mode only: U	info-up
	Info mode only: <Space>	info-next-page
	Info mode only: <Backspace>	info-previous-page
	Info mode only: [info-backward-node
	Info mode only:]	info-forward-node
	Info mode only: M	info-menu

Info mode only: 0, 1, 2, ...	info-nth-menu-item
Info mode only: F	info-follow-reference
Info mode only: <Tab>	info-next-reference
Info mode only: Shift-<Tab>	info-previous-reference
Info mode only: <Enter>	info-follow-nearest-reference
Info mode only: L	info-last
Info mode only: Q	info-quit
Info mode only: T	info-top
Info mode only: D	info-directory-node
Info mode only: S	info-search
Info mode only: I	info-index
Info mode only: <Comma>	info-index-next
Info mode only: >	info-last-node
Info mode only: G	info-goto
	info
	info-mode
	info-validate
	info-tagify

4.1.2 Web-based Epsilon Documentation

Epsilon's online manual is available in three formats:

- You can read the manual in an Epsilon buffer using Info mode by pressing F1 f. See page 39.
- Users running Microsoft Windows versions prior to Windows Vista can access the WinHelp version of the manual by pressing F1 w. See page 37 for more information.
- You can view the HTML version of the manual using a web browser by pressing F1 h.

To display the HTML manual, Epsilon starts a documentation server program. This is named `lhelp.exe` (or `lhelpd` in Unix). The documentation server runs in the background, hiding itself from view, and your web browser communicates with it on a special "port", as if it were a web server.

The documentation server must be running in order to serve documentation, so a bookmark to a page in the documentation will only work if the documentation server is running. You can press F1 h in Epsilon to ensure it's running. To force an instance of the documentation server to exit, invoke it again with the `-q` flag.

If your browser is configured to use a proxy, you will typically need to tell it not to use proxy settings for addresses starting with 127.0.0.1 so that it may connect to the local documentation server.

Epsilon for Unix uses a shell script named `goto_url` to run a browser. You can edit it if you prefer a different browser. Epsilon will first look for a customized copy of `goto_url` in your `~/epsilon` directory. If there is none, it will search for and invoke any customized copy of `goto_url` it finds on your path. Failing that, it will use the standard copy installed in Epsilon's bin directory. Epsilon for Windows uses the system's default browser.

4.2 Moving Around

4.2.1 Simple Movement Commands

The most basic commands involve moving point around. Recall from page 24 that point refers to the place where editing happens.

The Ctrl-f command moves point forward one character, and Ctrl-b moves it back. Ctrl-a moves to the beginning of the line, and Ctrl-e moves to its end.

Ctrl-n and Ctrl-p move point to the next and previous lines, respectively. They will try to stay in the same column in the new line, but will never expand a line in order to maintain the column; instead they will move to the end of the line (but see below). The key Alt-< moves point before the first character in the buffer, and Alt-> moves point after the last character in the buffer.

You can use the arrow keys if you prefer: the <Right> key moves forward a character, <Left> moves back a character, <Down> moves down a line, and <Up> moves up a line. Most commands bound to keys on the numeric keypad also have bindings on some control or alt key for those who prefer not to use the keypad. Throughout the rest of this chapter, the explanatory text will only mention one of the bindings in such cases; the other bindings will appear in the summary at the end of each section.

By default, pressing <Right> at the end of the line moves to the start of the next line. When you press <Down> at the end of a 60-character line, and the next line only has 10 characters, Epsilon moves the cursor back to column 10. You can change this by setting the buffer-specific `virtual-space` variable (by default zero). If you set it to one, the <Up> and <Down> keys will stay in the same column, even if no text exists there. If you set it to two, in addition to <Up> and <Down>, the <Right> and <Left> keys will move into places where no text exists, always remaining on the same line of the buffer. Setting `virtual-space` to two only works correctly on lines longer than the window when Epsilon has been set to scroll long lines (the default), rather than wrapping them (see page 108). Some commands behave unexpectedly on wrapped lines when `virtual-space` is two.

When you move past the bottom or top of the screen using <Up> or <Down>, Epsilon scrolls the window by one line, so that point remains at the edge of the window. If you set the variable `scroll-at-end` (normally 1) to a positive number, Epsilon will scroll by that many lines when <Up> or <Down> would leave the window. Set the variable to 0 if you want Epsilon to instead center the current line in the window.

Summary:	Ctrl-a, Alt-<Left>	beginning-of-line
	Ctrl-e, Alt-<Right>	end-of-line
	Ctrl-n, <Down>	down-line
	Ctrl-p, <Up>	up-line
	Ctrl-f, <Right>	forward-character
	Ctrl-b, <Left>	backward-character
	Alt-<, Ctrl-<Home>	goto-beginning
	Alt->, Ctrl-<End>	goto-end

4.2.2 Moving in Larger Units

Words

Epsilon has several commands that operate on words. A word usually consists of a sequence of letters, numbers, and underscores. The Alt-f and Alt-b commands move forward and backward by words, and the Alt-d and Alt-<Backspace> commands kill forward and backward by words, respectively. Like all killing

commands, they save away what they erase (see page 59 for a discussion on the killing commands). Epsilon's word commands work by moving in the appropriate direction until they encounter a word edge.

The word commands use a regular expression to define the current notion of a word. They use the buffer-specific variable `word-pattern`. This allows different modes to have different notions of what constitutes a word. Most built-in modes, however, make `word-pattern` refer to the variable `default-word`, which you can modify. See page 66 for information on regular expressions, and page 159 for information on setting this variable.

You can set the `forward-word-to-start` variable nonzero if you want Epsilon to stop at the start of a word instead of at its end when moving forward.

Summary:	Alt-f, Ctrl-⟨Right⟩	forward-word
	Alt-b, Ctrl-⟨Left⟩	backward-word
	Alt-⟨Backspace⟩	backward-kill-word
	Alt-d	kill-word

Sentences

For sentences, Epsilon has the Alt-e and Alt-a keys, which move forward and backward by sentences, and the Alt-k key, which deletes forward to the end of the current sentence. A sentence ends with one of the characters period, !, or ?, followed by any number of the characters ", ',),], followed by a space or a newline. A sentence also ends at the end of a paragraph. The next section describes Epsilon's notion of a paragraph.

You can set the `sentence-end-double-space` variable to change Epsilon's notion of a sentence. The commands in this section will require two spaces to mark the end of a sentence, and paragraph filling commands will use two spaces as well. This older style also permits Epsilon to distinguish abbreviations (followed by one space) from the ends of sentences (two spaces), which it can't do by default.

Summary:	Alt-e	forward-sentence
	Alt-a	backward-sentence
	Alt-k	kill-sentence

Paragraphs

For paragraphs, the keys Alt-] and Alt-[move forward and back, and the key Alt-h puts point and mark around the current paragraph. Blank lines (containing only spaces and tabs) always separate paragraphs, and so does the form-feed character ^L.

You can control what Epsilon considers a paragraph using two variables.

If the buffer-specific variable `indents-separate-paragraphs` has a nonzero value, then a paragraph also begins with a nonblank line that starts with a tab or a space.

If the buffer-specific variable `tex-paragraphs` has a nonzero value, then Epsilon will not consider as part of a paragraph any sequence of lines that each start with at sign or period, if that sequence appears next to a blank line. And lines starting with `\begin` or `\end`, or with `%`, `\[`, `\]`, or `$$`, or ending with `\\`, will also delimit paragraphs.

Summary:	Alt-], Alt-⟨Down⟩	forward-paragraph
	Alt-[, Alt-⟨Up⟩	backward-paragraph
	Alt-h	mark-paragraph

Parenthetic Expressions

Epsilon has commands to deal with matching parentheses, square brackets, curly braces, and similar delimiters. We call a pair of these characters with text between them a *level*. You can use these level commands to manipulate expressions in many programming languages, such as Lisp, C, and Epsilon's own embedded programming language, EEL.

A level can contain other levels, and Epsilon won't get confused by the inner levels. For example, in the text "one (two (three) four) five" the string "(two (three) four)" constitutes a level. Epsilon recognizes that "(three)" also constitutes a level, and so avoids the mistake of perhaps calling "(two (three))" a level. In each level, the text inside the delimiters must contain matched pairs of that delimiter. In many modes, Epsilon knows to ignore delimiters inside strings or comments, when appropriate.

Epsilon typically recognizes the following pairs of enclosures: '(' and ')', '[' and ']', '{' and '}'. The command Ctrl-Alt-f moves forward to the end of the next level, by looking forward until it sees the start of a level, and moving to its end. The command Ctrl-Alt-b moves backward by looking back for the end of a level and going to its beginning. The Ctrl-Alt-k command kills the next level by moving over text like Ctrl-Alt-f and killing as it travels, and the Alt-⟨Del⟩ command moves backward like Ctrl-Alt-b and kills as it travels. A mode may define a different set of grouping characters, such as < and > for HTML mode.

The Alt-) key runs the find-delimiter command. Use it to temporarily display a matching delimiter. The command moves backward like Ctrl-Alt-b and pauses for a moment, showing the screen, then restores the screen as before. The pause normally lasts one half of a second, or one second if the command must temporarily reposition the window to show the matching delimiter. You can specify the number of hundredths of a second to pause by setting the variables near-pause and far-pause. Also, typing any key will immediately restore the original window context, without further pause.

The show-matching-delimiter command inserts the key that invoked it by calling normal-character and then invokes find-delimiter to show its match. The maybe-show-matching-delimiter command is similar, but only invokes find-delimiter if the Matchdelim variable is nonzero. In Fundamental mode, the ')', ']' and '}' keys run maybe-show-matching-delimiter.

In some modes, when the cursor is over or next to a delimiter, Epsilon will automatically seek out its matching delimiter and highlight them both. (The auto-show-adjacent-delimiter variable controls whether highlighting occurs when next to a delimiter, not on it.) See the descriptions of the individual modes for more information.

Summary:	Alt-)	find-delimiter
	Ctrl-Alt-f	forward-level
	Ctrl-Alt-b	backward-level
	Ctrl-Alt-k	kill-level
	Alt-⟨Del⟩	backward-kill-level
		show-matching-delimiter

4.2.3 Searching

Epsilon provides a set of flexible searching commands that incorporate *incremental search*. In the incremental-search command, Epsilon searches as you type the search string. Ctrl-s begins an incremental search forward, and Ctrl-r starts one in reverse. Any character that normally inserts itself into the buffer becomes part of the search string. In an incremental search, Ctrl-s and Ctrl-r find the next occurrence of the string in the forward and reverse directions, respectively. With an empty search string, Ctrl-s or Ctrl-r will either reverse the direction of the search, or bring in the previously used search string. (To retrieve older search strings, see page 30.)

You can use <Backspace> to remove characters from the search string, and enter control characters and *meta characters* (characters with the eighth bit set) in the search string by quoting them with Ctrl-q. (Type Ctrl-q Ctrl-j to search for a <Newline> character.) Use the Ctrl-g abort command to stop a long search in progress.

Typing <Enter> or <Esc> exits from an incremental search, makes Epsilon remember the search string, and leaves point at the match in the buffer.

While typing characters into the search string for incremental-search, a Ctrl-g quits and moves point back to the place the search started, without changing the default search string. During a failing search, however, Ctrl-g simply removes the part of the string that did not match.

If you type an editing key not mentioned in this section, Epsilon exits the incremental search, then executes the command bound to the key.

You can make Epsilon copy search text from the current buffer by typing Alt-<Down>. Epsilon will append the next word from the buffer to the current search string. This is especially convenient when you see a long variable name, and you want to search for other references to it. (It's similar to setting the mark and moving forward one word with Alt-f, then copying the text to a kill buffer and yanking it into the current search string.) Similarly, Alt-<PageDown> appends the next line from the current buffer to the search string. These two keys are actually available at almost any Epsilon prompt, though they're especially useful when searching. Alt-Ctrl-n and Alt-Ctrl-v are synonyms for Alt-<Down> and Alt-<PageDown>, respectively.

While Alt-<Down> and Alt-<PageDown> copy text from the buffer at point, using the word pulling keys F3, Ctrl-<Up> or Ctrl-<Down> copies text into the search string from other parts of the buffer; see page 101.

You can change how Epsilon interprets the search string by pressing certain keys when you type in the search string. Pressing the key a second time restores the original interpretation of the search string.

- Pressing Ctrl-c toggles the state of *case folding*. While case folding, Epsilon considers upper case and lower case the same when searching, so a search string of "Word" would match "word" and "WORD" as well.

Epsilon remembers the state of case folding for each buffer separately, using the buffer-specific variable `case-fold`. When you start to search, Epsilon sets its default for case folding based on that variable's value for the current buffer. Toggling case folding with Ctrl-c won't affect the default. Use the `toggle-case-fold` command to do this, or set the `case-fold` variable using the `set-variable` command described on page 159 to change the default for case folding.

- Pressing Ctrl-w toggles *word searching*. During word searching, Epsilon only looks for matches consisting of complete words. For instance, word searching for 'a' in this sentence finds only one match (the one in quotes), but five when not doing word searching. You can type multiple words separated by spaces, and Epsilon will recognize them no matter what whitespace characters separate them (for instance, if they're on successive lines).
- Pressing Ctrl-t makes Epsilon interpret the search string as a regular expression search pattern, as described on page 66. Another Ctrl-t turns off this interpretation. If the current search string denotes

an invalid regular expression, Epsilon displays “Bad R-E Search: <string>” instead of its usual message “R-E Search: <string>” (where <string> refers to the search string). (When word and regular expression modes are combined, the entire pattern must start and end on a word boundary. Use </word> to match word boundaries within it.)

- Pressing Ctrl-o toggles incremental searching. In an incremental search, most editing commands will exit the search, as described above. But you may want to edit the search string itself. If you turn off the “incremental” part of incremental search with the Ctrl-o key, Epsilon will let you use the normal editing keys to modify the search string.

In non-incremental mode, Epsilon won’t automatically search after you type each character, but you can tell it to find the next match by typing Ctrl-s or Ctrl-r (depending on the direction). This performs the search but leaves you in search mode, so you can find the next occurrence of the search string by typing Ctrl-s or Ctrl-r again. When you press <Enter> to exit from the search, Epsilon will search for the string you’ve entered, unless you’ve just searched with Ctrl-s or Ctrl-r. (In general, the <Enter> key causes a search if the cursor appears in the echo area. If, on the other hand, the cursor appears in a window showing you a successful search, then typing the <Enter> key simply stops the search.) A numeric argument of *n* to a non-incremental search will force Epsilon to find the *n*th occurrence of the indicated string.

Epsilon interprets the first character you type after starting a search with Ctrl-s or Ctrl-r a little differently. Normally, Ctrl-s starts an incremental search, with regular expression searching and word searching both disabled. If you type Ctrl-t or Ctrl-w to turn one of these modes on, Epsilon will also turn off incremental searching. Epsilon also pulls in a default search string differently if you do it immediately. It will always provide the search string from the last search, interpreting the string as it did for that search. If you retrieve a default search string at any other time, Epsilon will provide the last one consistent with the state of regular expression mode (in other words, the last regular expression pattern, if in regular expression mode, or the last non-regular-expression string otherwise).

There are other ways besides Ctrl-s or Ctrl-r to retrieve previous search strings. You can press Alt-(Up) or Ctrl-Alt-p to display a list of previous search patterns. Press <Enter> to select one. Or you can press Alt-g at a search prompt to retrieve the search string from your last search in the current buffer only. This can differ from the default search string you get when you use Ctrl-s or Ctrl-r, since those are not per-buffer.

The Ctrl-Alt-s and Ctrl-Alt-r commands function like Ctrl-s and Ctrl-r, but they start in regular-expression, non-incremental mode. You can also start a plain string search in non-incremental mode using the string-search and reverse-string-search commands. Some people like to bind these commands to Ctrl-s and Ctrl-r, respectively. Also see the `search-positions-at-start` variable.

Keep in mind that you can get from any type of search to any other type of search by typing the appropriate subcommands to a search. For example, if you meant to do a `regex-search` but instead typed Ctrl-s to do an incremental search, you could enter regex mode by typing Ctrl-t. Figure 4.1 summarizes the search subcommands.

When you’re at the last match of some text in a buffer, and tell incremental search to search again by pressing Ctrl-s, Epsilon displays “Failing” to indicate no more matches. If you press Ctrl-s once more, Epsilon will wrap to the beginning of the buffer and continue searching from there. It will display “Wrapped” to indicate it’s done this. If you keep on search, eventually you’ll pass your starting point again; then Epsilon will display “Overwrapped” to indicate that it’s showing you a match you’ve already seen. A reverse search works similarly; Epsilon will wrap to the end of the buffer when you keep searching after a search has failed. (You can set the `search-wraps` variable to zero to disable wrapping.)

In some modes like Info mode, where a buffer displays a single part of some larger collection of text, pressing Ctrl-s at a failing search results in a continued search, instead of wrapping. Epsilon displays “Continued” to indicate (in the case of Info mode) that it’s searching through other nodes.

- Ctrl-s or Ctrl-r** Switch to a new direction, or find the next occurrence in the same direction, or pull in the previous search string.
- normal key** Add that character to the search string.
- ⟨Backspace⟩** Remove the last character from the search string.
- Ctrl-g** Stop a running search, or (in incremental mode) delete characters until the search succeeds, or abort the search, returning to the starting point.
- Ctrl-o** Toggle incremental searching.
- Ctrl-t** Toggle regular expression searching.
- Ctrl-w** Toggle word searching. Matches must consist of complete words.
- Ctrl-c** Toggle case folding.
- ⟨Enter⟩** Exit the search.
- Ctrl-d or ⟨Del⟩** Delete the current match and exit the search (but see the `search-delete-match` variable).
- Ctrl-q** Quote the following key, entering it into the search string even if it would normally run a command.
- help key** Show the list of search subcommands.
- other keys** If in incremental mode, exit the search, then execute the key normally. If not incremental mode, edit the search string.

Figure 4.1: The search subcommands work in all search and replace commands.

The `forward-search-again` and `reverse-search-again` commands search forward and backward (respectively) for the last-searched-for search string, without prompting. The `search-again` command searches in the same direction as before for the same search string.

The `search-region` command restricts searching to the current region, which will be highlighted during the search command.

If you highlight a region before searching, Epsilon uses it as an initial search string if it's not very long. Set the `search-in-region` variable to make Epsilon instead restrict matches it finds to the highlighted region, like the `search-region` command. Also see the `search-defaults-from` variable.

You can change the function of most keys in Epsilon by rebinding them (see page 157). But Epsilon doesn't implement the searching command keys listed above with the normal binding mechanism. The EEL code for searching refers directly to the keys `Ctrl-c`, `Ctrl-w`, `Ctrl-t`, `Ctrl-o`, `Ctrl-q`, `⟨Enter⟩`, and `⟨Esc⟩`, so to change the function of these keys within searching you must modify the EEL code in the file `search.e`. Epsilon looks at your current bindings to determine which keys to use as the help key and backspace key. It looks at the `abort_key` variable to determine what to use as your abort key, instead of `Ctrl-g`. (See page 106.) Epsilon always recognizes `Ctrl-s` and `Ctrl-r` as direction keys, but you can set two variables `fwd-search-key` and `rev-search-key` to key codes. These will then act as “synonyms” to `Ctrl-s` and `Ctrl-r`, respectively.

When you select a searching command from the menu or tool bar (rather than via a command's keyboard binding), Epsilon for Windows runs the `dialog-search` or `dialog-reverse-search` command, to display a search dialog.

Most of the keys described above also work in dialog-based searching. However, dialog searching is

never incremental, so Ctrl-o doesn't toggle incremental searching in a dialog. And Ctrl-q doesn't quote the following character, because dialog searching doesn't support directly entering special characters.

To match special characters in dialog-based searching, you can enable regular expression searching, and then enter them using syntax like <Tab> or <#13>. See page 67. In replacement text, add a # first, as in #<Newline> or #<#13>. See page 75.

Summary:	Ctrl-s	incremental-search
	Ctrl-r	reverse-incremental-search
	Ctrl-Alt-s	regex-search
	Ctrl-Alt-r	reverse-regex-search
		string-search
		reverse-string-search
		search-again
		forward-search-again
		reverse-search-again
		search-region
		dialog-search
		dialog-reverse-search
		toggle-case-fold

Searching Multiple Files

Epsilon provides a convenient `grep` command that lets you search a set of files. The command prompts you for a search string (all of the search options described above apply) and for a file pattern. By default, the `grep` interprets the search string as a regular expression (see page 66). To toggle regular expression mode, press Ctrl-t at any time while typing the search string. The command then scans the indicated files, puts a list of matching lines in the `grep` buffer, then displays the `grep` buffer in the current window. Each line indicates the file it came from.

With a numeric argument, this command searches through buffers instead of files. Instead of prompting for a file name pattern, Epsilon prompts for a buffer name pattern, and only operates on those buffers whose names match that pattern. Buffer name patterns use a simplified file name pattern syntax: * matches zero or more characters, ? matches any single character, and character classes like [a-z] may be used too. The `buffer-grep` command is an equivalent way to search buffers, handy if you want to bind it to its own key, and the command `this-buffer-grep` provides a similar shortcut, omitting the prompt for which buffer.

When `grep` prompts for a file pattern, it shows you the last file pattern you searched inside square brackets. You can press <Enter> to conveniently search through the same files again. (See the `grep-default-directory` variable to control how Epsilon interprets this default pattern when the current directory has changed.)

By default file patterns you type are interpreted relative to the current buffer's file; see `grep-prompt-with-buffer-directory` to change this. To repeat a file pattern from before, press Alt-(Up) or Ctrl-Alt-p. (See page 30 for details.) You can use extended file patterns to search in multiple directories; see page 138.

Epsilon skips over any file with an extension listed in `grep-ignore-file-extensions`; by default some binary file types are excluded. It skips files bigger than 2 GB, the default setting of the `grep-ignore-file-max-size` variable. It also skips over files matched by the `grep-ignore-file-pattern` or `grep-ignore-file-basename` variables (the latter matched against just

the base name of the file, not its path, the former matched against the entire file name). The `grep-ignore-file-types` variable makes `grep` skip over files that refer to devices, named pipes, or other sorts of special files. You can set the `use-grep-ignore-file-variables` variable to zero temporarily to have Epsilon ignore all these variables and search every matching file, or use the `select-and-grep` command, which lets you define custom rule sets for `grep`. By default, there's a rule set that searches all files, ignoring all the restriction rules in this paragraph.

In a `grep` buffer, you can move around by using the normal movement commands. Most alphabetic keys run special `grep` commands. The 'N' and 'P' keys move to the next and previous matches. The `Alt-n` and `Alt-p` keys move to the next and previous files. `Alt-]` and `Alt-[` move to the next and previous searches.

You can easily go from the `grep` buffer to the corresponding locations in the original files. To do this, simply position point on the copy of the line, then press `<Space>`, `<Enter>`, or 'E'. The file appears in the current window, with point positioned at the beginning of the matching line. Typing '1' brings up the file in a window that occupies the entire screen. Typing '2' splits the window horizontally, then brings up the file in the lower window. Typing '5' splits the window vertically, then brings up the file. Typing the letter 'O' shows the file in the next window on the screen, without splitting windows any further. Typing 'Z' runs the `zoom-window` command, then brings up the file.

When Epsilon wants to search a particular file as a result of a `grep` command, it first scans the buffers to see if one of them contains the given file. If so, it uses that buffer. If the file doesn't appear in any buffer, Epsilon reads the file into a temporary buffer, does the search, then discards the buffer.

If you want Epsilon to always keep the files around in such cases, set the variable `grep-keeps-files` to a nonzero value. In that case, `grep` will simply use the `find-file` command to get any file it needs to search.

By default, each invocation of `grep` appends its results to the `grep` buffer. If you set the variable `grep-empties-buffer` to a nonzero value, `grep` will clear the `grep` buffer at the start of each invocation. Also see the `grep-show-absolute-path` variable to control the format of file names in the `grep` buffer, and the `wrap-grep` variable to control whether grepping sets the current window to wrap long lines.

You can move from match to match without returning to the `grep` buffer. The `Ctrl-x Ctrl-n` command moves directly to the next match. It does the same thing as switching to the `grep` buffer, moving down one line, then pressing `<Space>` to select that match. Similarly, `Ctrl-x Ctrl-p` backs up to the previous match.

Actually, `Ctrl-x Ctrl-n` runs the `next-position` command. After a `grep` command, this command simply calls `next-match`, which moves to the next match as described above. If you run a compiler in a subprocess, however, `next-position` calls `next-error` instead, to move to the next compiler error message. If you use the `grep` command again, or press `<Space>` in the `grep` buffer to select a match, or run `next-match` explicitly, then `next-position` will again call `next-match` to move to the next match.

Similarly, `Ctrl-x Ctrl-p` actually runs `previous-position`, which calls either `previous-error` or `previous-match`, depending upon whether you last ran a compiler or searched across files.

Summary:	Alt-F7	<code>grep</code>
		<code>buffer-grep</code>
		<code>this-buffer-grep</code>
		<code>select-and-grep</code>
	Ctrl-x Ctrl-n	<code>next-position</code>
	Ctrl-x Ctrl-p	<code>previous-position</code>
		<code>next-match</code>
		<code>previous-match</code>

4.2.4 Bookmarks

Epsilon's bookmark commands let you store the current editing position, so that you can easily return to it later. To drop a bookmark at point, use the Alt-/ key. For each bookmark, Epsilon remembers the buffer and the place within that buffer. Later, when you want to jump to that place, press Alt-j. Epsilon remembers the last 10 bookmarks that you set with Alt-/. To cycle through the last 10 bookmarks, you can press Alt-j and keep pressing it until you arrive at the desired bookmark.

The tagging commands set a bookmark automatically before they move to a new location. You can set the `set-bookmark-automatically` variable to 1 if you want searching commands to do the same.

You can set a named bookmark with the Ctrl-x / key. The command prompts you for a letter, then associates the current buffer and position with that letter. To jump to a named bookmark, use the Ctrl-x j key. It prompts you for the letter, then jumps to that bookmark.

Instead of a letter, you can specify a digit (0 to 9). In that case, the number refers to one of the temporary bookmarks that you set with the Alt-/ key. Zero refers to the last temporary bookmark, 1 to the one before that, and so on.

Whenever one of these commands asks you to specify a character for a bookmark, you can get a list by pressing '?'. Epsilon then pops up a list of the bookmarks you've defined, along with a copy of the line that contains the bookmark. You can simply move to one of the lines and press <Enter> to select that bookmark. In a list of bookmarks, press D to delete the highlighted bookmark.

The command `list-bookmarks` works like the Ctrl-x j key, but automatically pops up the list of bookmarks to choose from. If you like, you can bind it to Ctrl-x j to get that behavior.

Summary:	Alt-/	<code>set-bookmark</code>
	Alt-j	<code>jump-to-last-bookmark</code>
	Ctrl-x /	<code>set-named-bookmark</code>
	Ctrl-x j	<code>jump-to-named-bookmark</code>
		<code>list-bookmarks</code>

4.2.5 Tags

Epsilon provides a facility to remember which file defines a particular subroutine or procedure. This can come in handy if your program consists of several source files. Epsilon can remember this kind of information for you by using "tags". A tag instructs Epsilon to look for a particular function at a certain position in a certain file.

The `goto-tag` command on Ctrl-x <Period> prompts for the name of a function and jumps immediately to the definition of the routine. You can use completion (see page 28) while typing the tag name, or press '?' to select from a list of tags. (Epsilon also shows the defining file of each tag.)

If you don't give a name, `goto-tag` goes to the next tag with the same name as the last tag you gave it. If the same tag occurs several times (for example, if you tag several separate files that each define a `main()` function), use this to get to the other tag references, or press '?' after typing the tag name to select the correct file from a list. If you give `goto-tag` a nonzero numeric argument, it goes to the next tag without even asking for a name. When there are several instances of a single tag, you can also use Ctrl-(NumPlus) and Ctrl-(NumMinus) to move among them.

The `plug-tag` command on Ctrl-x <Comma> first retrieves the routine name adjacent to or to the right of point, then jumps to that routine's definition.

If the file containing the definition appears in a window already, Epsilon will change to that window. Otherwise, Epsilon uses the find-file command to read the file into a buffer and displays it in the current window. Then Epsilon jumps to the definition, positioning its first line near the top of the window. You can set the window line to receive the first line of the definition via the `show-tag-line` variable. It says how many lines down the definition should go.

You can tell Epsilon to display the definition in a particular window, instead of letting Epsilon decide, by running `goto-tag` or `pluck-tag` with a numeric prefix argument of zero. Then these commands will prompt for a key to indicate the window. Press an arrow key to display the definition in the next window in that direction. Press `n` or `p` to display the definition in the next or previous window in the window order. Type the period character `.` to force the definition to appear in the current window. Press `2` or `5` to split the current window horizontally or vertically, respectively, and display the definition in the new window, or `1` to delete all windows but the current one, or `z` to run the `zoom-window` command first.

Before Epsilon moves to the tag, it sets a temporary bookmark at your old position, just like the `set-bookmark` command on `Alt-/'`. After `goto-tag` or `pluck-tag`, press `Alt-j` or `Ctrl-(NumStar)` to move back to your previous position.

Normally, you have to tell Epsilon beforehand which files to look in. The `tag-files` command on `Ctrl-x Alt-(Period)` prompts for a file name or file pattern such as `*.c` and makes a tag for each routine in the file. It knows how to recognize routines in C, C++, Java, Perl, Visual Basic, Python, PHP and many other languages. (Using EEL, you can teach Epsilon to tag additional languages. See page 519.) If you tag a previously tagged file, the new tags replace all the old tags for that file. You can use extended file patterns to tag files in multiple directories; see page 138. To easily tag just the current file, press `Alt-g` at the prompt. When Epsilon can't find a tag, it tries retagging the current file before giving up; that means if your program is confined to one file, you don't have to tag it first. Set `tag-ask-before-retagging` nonzero if you want Epsilon to ask first.

In Perl, PHP, Visual Basic, and Python, Epsilon tags subroutine definitions. In C, C++, Java, EEL and other C-like languages, `tag-files` normally tags subroutine and variable definitions, typedef definitions, structure and union member and tag definitions, enum constants, and `#define` constants. But it doesn't tag declarations (variables that use `extern`, function declarations without a body). With a numeric prefix argument, Epsilon includes these too. (Typically you'd do this for header files when you don't have source code for the function definitions—system files and library files, for instance.)

You can also set up `tag-files` to include declarations by default, by setting the `tag-declarations` variable. If zero (the default), `tag-files` only tags definitions. If one, Epsilon tags function declarations as well. If two, Epsilon tags variable declarations (which use the `extern` keyword). If three, Epsilon tags both types of declarations. Using a prefix argument with `tag-files` temporarily sets `tag-declarations` to three, so it tags everything it can. You can also set the `tag-which-items` variable to make tagging skip certain types of items, such as structure tag names or `#define` constants. Set `tag-c-preprocessor-skip-pat` to make Epsilon skip certain `#if` blocks when tagging C mode files.

Set `tag-case-sensitive` nonzero if you want tagging to consider `MAIN`, `Main` and `main` to be distinct tags. By default, typing `"main"` will find any of these.

Epsilon can maintain separate groups of tags, each in a separate file. The `select-tag-file` command on `Ctrl-x Alt-(Comma)` prompts for the name of a tag file, and uses that file for tag definitions.

When Epsilon needs to find a tag file, it searches for a file in the current directory, then in its parent directory, then in that directory's parent, and so forth, until it reaches the root directory or finds a file `"default.tag"`. If Epsilon finds no file with that name, it creates a new tag file in the current directory. To force Epsilon to create a new tag file in the current directory, even if a tag file exists in a parent directory, use the `select-tag-file` command. Once Epsilon loads a tag file, it continues to use that tag file until you use the `select-tag-file` command to select a new one, or delete the buffer named `"-tags"` (causing Epsilon to search again the next time you use a tagging command).

You can set the variable `initial-tag-file` to a relative pathname like “myfile.tag”, if you want Epsilon to search for that file, or you can set it to an absolute pathname if you want Epsilon to use the same tag file no matter which directory you use.

The tag system can also use .bsc files from Microsoft Visual Studio 5.0 and later. To use .bsc files, you must set your compiler to generate them, then use the `Alt-x configure-epsilon` command to download and install the DLL file that matches your compiler version. See page 52 for details. Finally, use the `select-tag-file` command on `Ctrl-x Alt-(Comma)` to select your .bsc file.

When Epsilon uses a .bsc file, the commands `tag-files`, `retag-files`, `clear-tags`, `sort-tags`, and the variables `tag-case-sensitive`, `tag-relative`, `want-sorted-tags`, and `tag-by-text` do not apply.

The `retag-files` command makes Epsilon rescan all the files represented in the current tag file and generate a new set of tags for each, replacing any prior tags. The `clear-tags` command makes Epsilon forget about all the tags in the current tag file. See the `tag-options` variable if you want the `tag-files` command to clear old tags automatically. The `untag-files` command displays a list of all files mentioned in the current tag file; you can edit the list by deleting any file names that shouldn’t be included, and when you press `Ctrl-x Ctrl-z`, Epsilon will forget all tags that refer to the file names you deleted.

When Epsilon records a tag, it stores the character position and the text of the line at the tag position. If the tag doesn’t appear at the remembered character offset, Epsilon searches for the defining line. And if that doesn’t work (perhaps because its defining line has changed) Epsilon retags the file and tries again. This means that once you tag a file, it should rarely prove necessary to retag it, even if you edit the file. To save space in the tag file, you can have Epsilon record only the character offset, by setting the variable `tag-by-text` to zero. Because this makes Epsilon’s tagging mechanism faster, it’s a good idea to turn off `tag-by-text` before tagging any very large set of files that rarely changes.

By default, Epsilon sorts the tag list whenever it needs to display a list of tag names for you to choose from. Although Epsilon tries to minimize the time taken to sort this list, you may find it objectionable if you have many tags. Instead, you can set the `want-sorted-tags` variable to 0, and sort the tags manually, whenever you want, using the `sort-tags` command. You can also tell Epsilon not to automatically save its tag file by setting the `auto-save-tags` variable to zero.

Epsilon normally stores file names in its tag file in relative format, when possible. This means if you rename or copy a directory that contains some source files and a tag file for them, the tag file will still work fine. If you set the variable `tag-relative` to 0, Epsilon will record each file name with an absolute pathname instead.

Summary:	<code>Ctrl-x (Period)</code>	<code>goto-tag</code>
	<code>Ctrl-x (Comma)</code>	<code>pluck-tag</code>
	<code>Ctrl-x Alt-(Period)</code>	<code>tag-files</code>
	<code>Ctrl-x Alt-(Comma)</code>	<code>select-tag-file</code>
	<code>Ctrl-(NumPlus)</code>	<code>next-tag</code>
	<code>Ctrl-(NumMinus)</code>	<code>previous-tag</code>
		<code>retag-files</code>
		<code>clear-tags</code>
		<code>untag-files</code>
		<code>sort-tags</code>

4.2.6 Source Code Browsing Interface

Epsilon can access source code browsing data generated by some Microsoft compilers. Compilers in Visual Studio versions 5.0 and later support this feature, though after Visual Studio 7.0 it is necessary to explicitly

enable it.

To set this up, first you must make sure your compiler generates such data, in the form of a .bsc file. If you build from the command line, compile with the /FR or /Fr flag to generate .sbr files, then use the bscmake utility to combine the .sbr files into a .bsc file. Or using the Visual Studio IDE, ensure the “Generate browse info” option (Project/Settings, on the C/C++ tab in the General category) and the “Build browse info file” option (Project/Settings, on the Browse Info tab) are both enabled. (The IDE options for this may vary by Visual Studio version.)

Next, set up Epsilon to use the generated browser file. To do this, run the Alt-x configure-epsilon command and select the option to install source code browser support. This retrieves the appropriate DLL and installs it.

You can use the browser database only for source code browsing, or you can tell Epsilon to use it for tagging as well, instead of using its own tagging methods. To have Epsilon use the same browser database file for both purposes, use the select-tag-file command on Ctrl-x Alt-⟨Comma⟩ to select your .bsc file. To use Epsilon’s built-in tagging, and utilize the browser database only for source code browsing, select your .bsc file with the select-browse-file command, which sets the browser-file variable.

Once you’ve set up source code browsing, press Ctrl-⟨NumSlash⟩ (using the / key on the numeric keypad) to run the browse-symbol command. It will prompt for the name of a symbol (the name of a function, variable, macro, class, or similar), using the symbol at point as the default. Then it will set a temporary bookmark at your old position, just like the set-bookmark command on Alt-/. (After using browse-symbol to navigate to a different part of your code, you can use Alt-j or Ctrl-⟨NumStar⟩ to move back to your original location.) Finally, it builds a #symbols# buffer showing all available information on the symbol.

The #symbols# buffer contains a header section, followed by one section for each distinct use of the symbol. For instance, if you use the name “cost” for a function, and also use it elsewhere as a local variable name, and as a structure name somewhere else, there will be three sections, one for each use.

```
Browser File: c:\Project\project.bsc
Symbol: qsort
Filter all but: Var Func Macro Type Class
Filter all Uses/UsedBy but: Var Func Macro Type Class
```

```
qsort (public function) is defined at:
qsort (public function) is used at:
- C:\Program Files\Microsoft Visual Studio\VC98\include\stdlib.h(302):
-- _CRTIMP void __cdecl qsort(void *, size_t, size_t, int (__cdecl *)
- prep_env.c(79): qsort(order, cnt, sizeof(char *), env_compare);
- token.cpp(174): qsort(1e, sizeX + sizeY, sizeof(line_entry), hash_cmp);
qsort (public function) is used by:
- make_proc_env (public function)
- tokenize(int *,int *) (static function)
```

The header section displays the name of the .bsc file used to generate the listing and the symbol being displayed. It also shows the current filters, which may be used to hide certain uses of a symbol.

Next you will see a list of those lines in your source code that define the specified symbol, followed by those lines that reference it. In the example above, qsort(), a library function, isn’t defined within the project source code, so its “is defined at” section is empty. It’s defined in a standard header file, and called from two places in the project source code. You can position to any of these source code lines and press ⟨Enter⟩, or double-click the line, and Epsilon will go to the corresponding source file and line.

In the following section, you will see a list of functions that use the `qsort()` function. You can look up any one of these symbol names with `<Enter>` or double-clicking, and Epsilon will display the symbol listing for that symbol, replacing the current listing. Afterwards, press the `L` key to return to viewing the original symbol. Repeated presses go to earlier symbols. With a numeric argument, the `L` key displays a list of recently-viewed symbols; you can select one and have it displayed again.

If the symbol has a definition within the current project, the next section will show the functions and variables it uses in its definition.

You can set filters, as shown in the header section, to skip over certain kinds of definitions and uses. For instance, if `qsort` were the name of a macro as well as a function, you could use the first filter to see only the macro uses by pressing `f`. The second filter controls which symbols appear in the uses/used-by section; press `b` to set it. You can also set the filters by pressing `<Enter>` while on the corresponding “Filter:” line in the browser buffer. These set the `browser-filter` and `browser-filter-usedby` variables.

The `browser-options` variable lets you omit some of the above sections, or simplify the data shown in other ways, to make browsing quicker. The `browse-current-symbol` command is a variation on `browse-symbol` that doesn’t prompt for a symbol name, but uses the name at point without prompting.

Summary:	<code>Ctrl-(NumSlash)</code>	<code>browse-symbol</code>
		<code>browse-current-symbol</code>
	Browse mode only: <code>f</code>	<code>browse-set-filter</code>
	Browse mode only: <code>b</code>	<code>browse-set-usedby-filter</code>
		<code>select-browse-file</code>

4.2.7 Comparing Two Buffers

The `compare-windows` command on `Ctrl-F2` finds differences between the contents of the current buffer and that displayed in the next window on the screen. If called while in the last window, it compares that window with the first window. The comparison begins at point in each window. Epsilon finds the first difference between the buffers and moves the point to just before the differing characters, or to the ends of the buffers if it finds no difference. It then displays a message in the echo area reporting whether or not it found a difference.

If you invoke `compare-windows` again immediately after it has found a difference, the command will try to resynchronize the windows by moving forward in each window until it finds a match of at least `resynch-match-chars` characters. It doesn’t necessarily move each window by the same amount, but instead finds a match that minimizes the movement in the window that it moves the most. It then reports the number of characters in each window it skipped past.

Normally `compare-windows` treats one run of space and tab characters the same as any other run, so it skips over differences in horizontal whitespace. You can set the `compare-windows-ignores-space` variable to change this.

The `diff` command also compares the buffers in two windows, but it will compare and resynchronize over and over from the beginning to the end of each buffer, producing a report that lists all differences between the two buffers. It operates line-by-line rather than character-by-character.

When resynchronizing, `diff` believes it has found another match when `diff-match-lines` lines in a row match, and gives up if it cannot find a match within `diff-mismatch-lines` lines. By default, `diff` resynchronizes when it encounters three lines in a row that match. Normally Epsilon uses a smarter algorithm that’s better at finding a minimum set of differences. With this algorithm, `diff-mismatch-lines` isn’t used. But because this algorithm becomes very slow when buffers are large,

it's only used when at least one of the buffers contains fewer than `diff-precise-limit` bytes (by default 4 MB).

The `diff` command reports each difference with a summary line and then the text of the differing lines. The summary line consists of two line number ranges with a letter between them indicating the type of change: 'a' indicates lines to add to the first buffer to match the second, 'd' indicates lines to delete, and 'c' indicates lines to change. For example, a summary line in the `diff` listing of "20,30c23,29" means to remove lines 20 through 30 from the first buffer and replace them with a copy of lines 23 through 29 from the second buffer. "11a12" means that adding line 12 from the second buffer right after line 11 in the first buffer would make them identical. "11,13d10" means that deleting lines 11, 12 and 13 from the first buffer (which would appear just after line 10 in the second) would make them identical.

After each summary line, `diff` puts the lines to which the summary refers. The `diff` command prefixes lines to delete from the first buffer by "<" and lines to add by ">".

The `visual-diff` command is like `diff` but uses colors to show differences. It constructs a new buffer that contains all the lines of the two buffers. Lines from the first buffer that don't appear in the second are displayed with a red background. Lines in the second buffer that don't appear in the first have a yellow background. Lines that are the same in both buffers are colored normally.

This command also does character-by-character highlighting for each group of changed lines. Instead of simply indicating that one group of lines was replaced by another, it shows which portions of the lines changed and which did not, by omitting the red or yellow background from those characters. You can set the variables `diff-match-characters` and `diff-match-characters-limit` to alter or turn off this behavior.

In a `visual-diff` buffer, the keys `Alt-(Down)` and `Alt-]` move to the start of the next changed or common section. The keys `Alt-(Up)` and `Alt-[` move to the previous change or common section. These keys work in `diff` buffers too, moving to the next or previous block of differences.

In both `visual-diff` and `diff` buffers, you can press "a" or "b" to visit either of the two files being compared. Press "c" to visit both in a pair of windows, one below the other, or "s" to visit both in side-by-side windows. Press `Shift-A` to copy file A to file B, or `Shift-B` to copy file B to file A. You will be prompted to confirm overwriting the existing file. Overwriting a newer file with an older one will mention that in the prompt.

The `merge-diff` command is another variation on `diff` that's useful with buffers in C mode. It marks differences by surrounding them with `#ifdef` preprocessor lines, first prompting for the `#ifdef` variable name to use. The resulting buffer receives the mode and settings of the first of the original buffers. The marking is mechanical, and doesn't parse the text being marked off, so it may produce invalid code. For example, if an `#if` statement differs between the two buffers, the result will contain improperly nested `#if` statements like this:

```
#ifndef DIFFVAR
    #if DOSVERSION
#else // DIFFVAR
    #if MSDOSVERSION
#endif // DIFFVAR
```

Therefore, you should examine the output of `merge-diff` before trying to compile it.

The commands `diff`, `visual-diff`, and `merge-diff` all produce their output in a buffer named `#diff#`. With a numeric argument, they prompt instead for the destination buffer name.

The `compare-to-prior-version` command uses `visual-diff` to show the differences between the current version of a buffer and the one saved on disk. It can also compare the current version with the version prior to a certain number of editing operations. It prompts for the number of editing operations; entering zero

makes it compare the current buffer to the version of it on disk. The command can display its results using `merge-diff` or `diff` instead of `visual-diff`; see the `compare-to-prior-version-style` variable.

Summary: `Ctrl-F2`, `Ctrl-x c` `compare-windows`
 `diff`
 `visual-diff`
 `visual-diff-mode`
 `merge-diff`
 `compare-to-prior-version`

Keys	In Diff and Visual Diff modes
<code>Alt-]</code> , <code>Alt-(Down)</code>	<code>next-difference</code>
<code>Alt-[</code> , <code>Alt-(Up)</code>	<code>previous-difference</code>
<code>a</code> , <code>b</code>	View file <code>a</code> or file <code>b</code>
<code>c</code>	View both files
<code>s</code>	View both files side by side
<code>Shift-A</code>	Copy file <code>a</code> to file <code>b</code>
<code>Shift-B</code>	Copy file <code>b</code> to file <code>a</code>

4.2.8 Comparing Many Files

The `compare-files` command compares all files matching a wildcard file pattern to corresponding files in another directory, listing which files are identical, different, or missing.

The pattern may include `**` to match a directory hierarchy. For example, `dir1/**` matches all files in `dir1` or its subdirectories, with any level of nesting, while `dir1/**/*.h` matches only `.h` files somewhere within that directory hierarchy. If you use a directory name instead of a file pattern, Epsilon compares all files in the hierarchy (same as `**`). You can use Epsilon's extended file patterns here; see page 138. If you provide two extended file patterns, instead of a pattern and a directory, Epsilon compares only the files that match the two patterns.

By default, Epsilon compares files while ignoring differences in spaces, tabs, and newlines. That is, any run of one or more spaces, tabs, or newlines (in any combination) matches any other such run. You can use a different rule by running `compare-files` with a numeric prefix argument, which makes it ask several questions about how to compare.

First, it offers three options for a quicker comparison. Instead of comparing the text of files after reading them with appropriate line translation and Unicode decoding (the default), you can have it compare files byte-for-byte. (This is faster because files with different sizes on disk can be immediately noted as mismatched.) An even faster method is to simply compare two files' timestamps and sizes, without examining their contents at all. (Set the `compare-files-method` variable if you want a different option by default.) If one of the file patterns refers to a remote file (starts with `scp:`, for example), then Epsilon always compares file contents, not sizes or dates.

Next (unless it's comparing only timestamps and sizes), `compare-files` asks how to compare whitespace. You can choose the default rule where any run of spaces, tabs, or newlines matches any other run, a similar rule that just applies to space and tabs (not newlines), or an exact mode where text must match exactly. (Set the `compare-files-ignore-spaces` variable to change the default rule.) With a prefix argument, `compare-files` also asks for the name of the destination buffer.

The resulting listing will show pairs of files that match, pairs that don't, and files that are only present in one directory but not the other. Press "a" at any line in the result to visit that line's file in the first directory or pattern you specified, or "b" to visit the one in the second. Press "c" to open both files in a pair of windows, one above the other, or "s" to open both files in side-by-side windows.

You can press "v" or `<Enter>` on any line to see a visual-diff listing comparing those two files. Press "d" to show a traditional diff listing the differences. Press "l" for a diff listing of every differing file in the two hierarchies.

Press Shift-A to copy file A to file B, or Shift-B to copy file B to file A. You will be prompted to confirm overwriting an existing file. Overwriting a newer file with an older one will mention that in the prompt.

Summary:

compare-files

Keys	Compare Files mode
a, b	View file a or file b
c	View both files
s	View both files side by side
v, <code><Enter></code>	View a Visual Diff of these two files
d	View a diff listing of these two files
l	View a diff of all the differing files
Shift-A	Copy file a to file b
Shift-B	Copy file b to file a

4.3 Changing Text

4.3.1 Inserting and Deleting

When you type most alphabetic or numeric keys, they appear in the buffer before point. Typing one of these keys runs the command `normal-character`, which simply inserts the character that invoked it into the buffer.

When you type a character bound to the `normal-character` command, Epsilon inserts the character before point, so that the cursor moves forward as you type characters. Epsilon can also overwrite as you type. The `overwrite-mode` command, bound to the `<Ins>` key, toggles overwriting for the current buffer. If you give it a nonzero numeric argument (for example, by typing `Ctrl-u` before invoking the command, see page 27), it doesn't toggle overwriting, but turns it on. Similarly, a numeric argument of zero always turns off overwriting. Overwriting will occur for all characters except newline, and overwriting never occurs at the end of a line. In these cases the usual insertion will happen. The buffer-specific variable `over-mode` controls overwriting.

The `Ctrl-q` key inserts special characters, such as control characters, into the current buffer. It waits for you to type a character, then inserts it. This command ignores keys that don't represent characters, such as `<Home>` or `F3`. If you "quote" an Alt key in this way, Epsilon inserts the corresponding character with its high bit on. You can use this command for inserting characters like `Ctrl-z` that would normally execute a command when typed.

Sometimes you may want to insert a character whose numeric ASCII value you know, but you may not know which keystroke that character corresponds to. Epsilon provides an `insert-ascii` command on `Alt-#` for this purpose. It prompts you for a numeric value, then inserts the character with that value into the buffer. By default, the command interprets the value in base 10. You can specify a hexadecimal value by prefixing the characters "0x" to the number, or an octal value by prefixing the character "0o" to the number, or a

binary value by prefixing “0b”. For example, the numbers “87”, “0x57”, “0o127”, and “0b1010111” all refer to the same number, and they all would insert a “W” character if given to the `insert-ascii` command.

You can also use the name of a Unicode character inside angle brackets, like “<square root>”, with `Alt-#`. Press `?` to see a list of characters with their Unicode names. You can use completion on character names like this, and search in the list of names as usual.

In most environments you can type graphics characters by holding down the `Alt` key and typing the character’s value on the numeric keypad, but see the `alt-numpad-keys` variable. In some environments, `Epsilon` will automatically quote the character so that it’s inserted in the buffer and not interpreted as a command. (You may need to type a `Ctrl-q` first to quote the character in other environments.)

The `Ctrl-o` command inserts a newline after point (or, to put it another way, inserts a newline before point as usual, then backs up over it). Use this command to break a line when you want to insert new text in the middle, or to “open” up some space after point.

The `<Backspace>` key deletes the character before point, and the `` key deletes the character after point. In other words, `<Backspace>` deletes backwards, and `` deletes forwards. These commands usually do not save deleted characters in the kill ring (see the next section).

If you prefix these commands with a numeric argument of n , they will delete n characters instead of one. In that case, you can retrieve the deleted text from the kill ring with the `Ctrl-y` key (see the next section).

If `<Backspace>` or `` follows one of the kill commands, the deleted character becomes part of the text removed by the kill command. See the following section for information on the kill commands, and the `delete-options` variable to change this behavior.

The buffer-specific variable `delete-hacking-tabs` makes `<Backspace>` operate differently when deleting tabs or spaces. If 1, when `<Backspace>` deletes a tab, it first turns the tab into the number of spaces necessary to keep the cursor in the same column, then deletes one of the spaces. If 2, when `<Backspace>` deletes a space, it deletes additional spaces and tabs until it reaches the previous tab column. The first setting makes `<Backspace>` treat tabs more like spaces; the second makes it treat spaces more like tabs. Other bits in the variable limit the circumstances where `<Backspace>` does this; see the variable’s documentation for details.

The key `Alt-\ deletes spaces and tabs surrounding point.`

The `Ctrl-x Ctrl-o` command deletes empty lines adjacent to point, or lines that contain only spaces and tabs, turning two or more such blank lines into a single blank line. `Ctrl-x Ctrl-o` deletes a lone blank line. If you prefix a numeric argument of n , exactly n blank lines appear regardless of the number of blank lines present originally. With a highlighted region, the command does this at every sequence of one or more blank lines throughout the region.

Summary:	<code>Ctrl-q</code>	<code>quoted-insert</code>
	<code>Alt-#</code>	<code>insert-ascii</code>
	<code>Ctrl-o</code>	<code>open-line</code>
	<code>Ctrl-h</code> , <code><Backspace></code>	<code>backward-delete-character</code>
	<code>Ctrl-d</code> , <code></code>	<code>delete-character</code>
	<code>Alt-\< code=""><></code>	<code>delete-horizontal-space</code>
	<code>Ctrl-x Ctrl-o</code>	<code>delete-blank-lines</code>
	“normal keys”	<code>normal-character</code>
	<code><Ins></code>	<code>overwrite-mode</code>

4.3.2 The Region, the Mark, and Killing

Epsilon has many commands to erase characters from a buffer. Some of these commands save the erased characters away in a special group of buffers called *kill buffers*, and some do not.

In Epsilon's terminology, to *kill* means to delete text and save it away in a kill buffer, and to *delete* means simply to remove the text and not save it away. Any consecutive sequence of killing commands will produce a single block of saved text. The Ctrl-y command then yanks back the entire block of text, inserting it before point. (Even when Epsilon deletes text and doesn't save it, you can usually use the undo command to recover the text. See page 105.)

The Ctrl-k command kills to the end of the line, but does not remove the line separator. At the end of a line, though, it kills just the line separator. Thus, use two Ctrl-k's to completely remove a nonempty line. Give this command a numeric argument of n to kill exactly n lines, including the line separators. If you give the Ctrl-k command a negative numeric argument, $-n$, the command kills from the beginning of the previous n th line to point.

The kill-current-line command is an alternative to Ctrl-k. It kills the entire line in one step, including the line separator. The kill-to-end-of-line command kills the rest of the line. If point is at the end of the line, it does nothing. In Brief mode Epsilon uses these two commands in place of the kill-line command that's normally bound to Ctrl-k.

The commands to delete single characters will also save the characters if you give them a numeric argument (to delete that number of characters) or if they follow a command which itself kills text.

Several Epsilon commands operate on a *region* of text. To specify a region, move to either end of the region and press the Ctrl-@ key or the Ctrl-⟨Space⟩ key. This sets the *mark* to the current value of point. Then move point to the other end of the region. The text between the mark and point specifies the region.

When you set the mark with Ctrl-@, Epsilon turns on highlighting for the region. As you move point away from the mark, the region appears in a highlighted color. This allows you to see exactly what text a region-sensitive command would operate upon. To turn the highlighting off, type Ctrl-x Ctrl-h. The Ctrl-x Ctrl-h command toggles highlighting for the region. If you prefix a nonzero numeric argument, it turns highlighting on; a numeric argument of zero turns highlighting off.

You can also check the ends of the region with the Ctrl-x Ctrl-x command. This switches point and mark, to let you see the other end of the region. Most commands do not care whether point (or mark) refers to the beginning or the end of the region.

The mark-whole-buffer command on Ctrl-x h provides a quick way to set point and mark around the entire buffer.

Another way to select text is to hold down the Shift key and move around using the arrow keys, or the keys ⟨Home⟩, ⟨End⟩, ⟨PageUp⟩, or ⟨PageDown⟩. Epsilon will select the text you move through. The `shift-selects` variable controls this feature.

The Ctrl-w command kills the region, saving it in a kill buffer. The Ctrl-y command then yanks back the text you've just killed, whether by the Ctrl-w command or any other command that kills text. It sets the region around the yanked text, so you can kill it again with a Ctrl-w, perhaps after adjusting the region at either end. The Alt-w command works like Ctrl-w, except that it does not remove any text from the buffer; it simply copies the text between point and mark to a kill buffer.

Each time you issue a sequence of killing commands, Epsilon saves the entire block of deleted text as a unit in one of its kill buffers. The Ctrl-y command yanks back the last of these blocks. To access the other blocks of killed text, use the Alt-y command. It follows a Ctrl-y or Alt-y command, and replaces the retrieved text with an earlier block of killed text. Each time you press Alt-y, Epsilon substitutes a block from another kill buffer, cycling from most recent back through the oldest, and then around to the most recent again.

In normal use, you go to the place you want to insert the text and issue the Ctrl-y command. If this doesn't provide the right text, give the Alt-y command repeatedly until you see the text you want. If the text you want does not appear in any of the killed blocks, you can get rid of the block with Ctrl-w, since both Ctrl-y and Alt-y always place point and mark around the retrieved block.

By default, Epsilon provides ten kill buffers. You can set the variable `kill-buffers` if you want a different number of kill buffers. Setting this variable to a new value makes Epsilon throw away the contents of all the kill buffers the next time you execute a command that uses kill buffers.

The Alt-y command doesn't do anything if the region changed since the last Ctrl-y or Alt-y, so you can't lose text with a misplaced Alt-y. Neither of these commands changes the kill buffers themselves. The Alt-y command uses the undo facility, so if you've disabled undo, it won't work.

Epsilon can automatically reindent yanked text. By default it does this in C mode buffers. See page 79 for details. If you invoke Ctrl-y or Alt-y with a negative numeric prefix argument, by typing Alt-(Minus) Ctrl-y for example, the command won't reindent the yanked text, and will insert one copy. (Providing a positive numeric prefix argument makes Epsilon yank that many copies of the text. See page 155.)

Each time you issue a sequence of killing commands, all the killed text goes into one kill buffer. When a killing command follows a non-killing command, the text goes into a new kill buffer (assuming you haven't set up Epsilon to have only one kill buffer). You may sometimes want to append a new kill to the current kill buffer, rather than using the next kill buffer. That would let you yank all the text back at once. The Ctrl-Alt-w command makes an immediately following kill command append to a kill buffer instead of moving to a new one.

The Ctrl-y command can come in handy when entering text for another command. For example, suppose the current buffer contains a line with "report.txt" on it, and you now want to read in the file with that name. Simply kill the line with Ctrl-k and yank it back (so as not to change the buffer) then give the Ctrl-x Ctrl-f command (see page 118) to read in a file. When prompted for the file name, press Ctrl-y and the text "report.txt" appears as if you typed it yourself.

Pressing a self-inserting key like 'j' while text is highlighted normally deletes the highlighted selection, replacing it with the key. Pressing (Backspace) simply deletes the text. You can disable this behavior by setting the variable `typing-deletes-highlight` to zero. If you turn off this feature, you may also wish to set the variable `insert-default-response` to zero. At many prompts Epsilon will insert a highlighted default response before you start typing, if this variable is nonzero. You may also wish to set `typing-hides-highlight` if you've disabled `typing-deletes-highlight`, so pressing a self-inserting key turns off highlighting but doesn't delete anything.

You can use the `delete-region` command to delete the current region without saving it in a kill buffer; this is especially useful if you've set (Backspace) so it doesn't delete highlighted text.

In addition to the above commands which put the text into temporary kill buffers, Epsilon provides commands to make more permanent copies of text. The Ctrl-x x key copies the text in the region between point and mark to a permanent buffer. The command prompts you for a letter (or number), then associates the text with that letter. Thereafter, you can retrieve the text using the Ctrl-x y key. That command asks you for the letter, then inserts the corresponding text before point.

Summary:	Ctrl-@, Alt-@	set-mark
	Ctrl-x Ctrl-h	highlight-region
	Ctrl-x Ctrl-x	exchange-point-and-mark
	Ctrl-k	kill-line
	Ctrl-w	kill-region
	Alt-w	copy-region
	Ctrl-y	yank

Alt-y	yank-pop
Ctrl-Alt-w	append-next-kill
Ctrl-x x	copy-to-scratch
Ctrl-x y	insert-scratch
Ctrl-x h	mark-whole-buffer
	kill-current-line
	kill-to-end-of-line
	delete-region

4.3.3 Clipboard Access

In Windows, Epsilon's killing commands interact with the Windows clipboard. Similarly, Epsilon for Unix interacts with the X11 clipboard when running as an X program. You can kill text in Epsilon and paste it into another application, or copy text from an application and bring it into Epsilon with the yank command.

All commands that put text on the kill ring will also try to copy the text to the clipboard, if the variable `clipboard-access` is non-zero. You can copy the current region to the clipboard without putting it on the kill ring using the command `copy-to-clipboard`.

The yank command copies new text from the clipboard to the top of the kill ring. It does this only when the clipboard's contents have changed since the last time Epsilon accessed it, the clipboard contains text, and `clipboard-access` is non-zero. Epsilon looks at the size of the clipboard to determine if the text on it is new, so it may not always notice new text. You can force Epsilon to retrieve text from the clipboard by using the `insert-clipboard` command, which inserts the text on the clipboard at point in the current buffer.

If you prefer to have Epsilon ignore the clipboard except when you explicitly tell it otherwise, set `clipboard-access` to zero. You can still use the commands `copy-to-clipboard` and `insert-clipboard` to work with the clipboard. Unlike the transparent clipboard support provided by `clipboard-access`, these commands will report any errors that occur while trying to access the clipboard. If transparent clipboard support cannot access the clipboard for any reason, it won't report an error, but will simply ignore the clipboard. Epsilon also disables transparent clipboard support when running a keyboard macro, unless `clipboard-access` is 2.

When the buffer contains syntax-highlighted text, or other text with colors applied to it, you can have Epsilon construct an HTML version of the text that preserves the coloring. You can then use it in a web page, or further convert it using an external converter. Run the `copy-formatting-as-html` command to copy the current region to the clipboard in HTML format.

By default, when the Windows Console version of Epsilon puts characters on the clipboard, it lets Windows translate the characters from the OEM character set to Windows ANSI, so that national characters display correctly. Epsilon for Windows uses Windows ANSI like other Windows programs, so no translation is needed. See the description of the `clipboard-format` variable to change this.

When retrieving text from the clipboard, Epsilon sometimes performs conversions to similar but more basic characters. For instance, if you paste Unicode U+02DC SMALL TILDE, Epsilon may replace it with the ASCII tilde character `~`. See the `clipboard-convert-unicode` variable for details on Unicode conversions to and from the system clipboard.

On macOS systems, Epsilon converts from Mac line termination conventions when you paste text. The `clipboard-convert-mac-lines` variable controls this.

X11 has two different methods of transferring text between programs. The more modern method uses the clipboard, and explicit commands for cutting and pasting text. This is what Epsilon's commands for killing and yanking use.

But an older method uses the “primary selection” to transfer text. Traditionally, selecting text with a mouse sets the text as the primary selection, and the middle mouse button pastes that text into another program.

In Epsilon the middle mouse button provides panning by default, but when you hold down Shift, it inserts the primary selection instead. You can set the `mouse-center-yanks` variable to make the middle mouse button always insert. Or you can use the `yank-x-selection` command to yank X’s primary selection explicitly. Set the `mouse-selection-copies` variable to make selecting text with the mouse set the primary selection. This also puts the text into one of Epsilon’s kill buffers.

If you mostly use programs that follow the older X11 convention, you can set Epsilon to do so as well. Set the `clipboard-format` variable to 1. Then Epsilon’s cutting and pasting commands will use the primary selection instead of the clipboard selection.

Summary:

`copy-to-clipboard`
`insert-clipboard`
`copy-formatting-as-html`
`yank-x-selection`

4.3.4 Rectangle Commands

Epsilon regions actually come in four distinct types. Each type has a corresponding Epsilon command that begins defining a region of that type.

Region Type	Command
Normal	<code>mark-normal-region</code>
Line	<code>mark-line-region</code>
Inclusive	<code>mark-inclusive-region</code>
Rectangular	<code>mark-rectangle</code>

The commands are otherwise very similar. Each command starts defining a region of the specified type, setting the mark equal to point and turning on highlighting. If Epsilon is already highlighting a region of a different type, these commands change the type. If Epsilon is already highlighting a region of the same type, these commands start defining a new region by setting mark to point again. (You can set the variable `mark-unhighlights` to make the commands turn off the highlighting and leave the mark alone in this case.)

The `mark-normal-region` command defines the same kind of region as the `set-mark` command described in section 4.3.2. (The commands differ in that `set-mark` always begins defining a new region, even if another type of region is highlighted on the screen. The `mark-normal-region` command converts the old region, as described above.)

A line region always contains entire lines of text. It consists of the line containing point, the line containing mark, and all lines between the two.

An inclusive region is very similar to a normal region, but an inclusive region contains one additional character at the end of the region. A normal region contains all characters between point and mark, if you think of point and mark as being positioned between characters. But if you think of point and mark as character positions, then an inclusive region contains the character at point, the character at the mark, and all characters between the two. An inclusive region always contains at least one character (unless point and mark are both at the end of the buffer).

A rectangular region consists of all columns between those of point and mark, on all lines in the buffer between those of point and mark. The mark-rectangle command on Ctrl-x # begins defining a rectangular region. In a rectangular region, point can specify any of the four corners of this rectangle.

Some commands operate differently when the current region is rectangular. Killing a rectangular region by pressing the Ctrl-w key runs the command kill-rectangle. It saves the current rectangle in a kill buffer, and replaces the rectangle with spaces, so as not to shift any text that appears to the right of the rectangle. To remove the rectangle and the space it occupied, press Ctrl-u Ctrl-w. This shifts columns of text that followed the rectangle to the left. (Also see the kill-rectangle-removes variable.)

The Alt-w key runs the command copy-rectangle. It also saves the current rectangle, but doesn't modify the buffer. (Actually, it may insert spaces at the ends of lines, or convert tabs to spaces, if that's necessary to reach the starting or ending column on one of the lines in the region. But the buffer won't look any different as a result of these changes. Most rectangle commands do this.)

The Ctrl-Alt-w key runs the command delete-rectangle. It removes the current rectangle, shifting any text after it to the left. It doesn't save the rectangle.

When you use the Ctrl-y key to yank a kill buffer that contains a rectangle, Epsilon inserts the last killed rectangle into the buffer at the current column, on the current and successive lines. It shifts existing text to the right. If you've enabled overwrite mode, however, the rectangle replaces any existing text in those columns. See the yank-rectangle-to-corner variable to set how Epsilon positions point and mark around the yanked rectangle. You can use the Alt-y key to cycle through previous kills as usual.

When yanking line regions, the yank-line-retains-position variable serves a similar purpose, influencing where Epsilon positions the cursor.

The width of a tab character depends upon the column it occurs in. For this reason, if you use the rectangle commands to kill or copy text containing tabs, and you move the tabs to a different column, text after the tabs may shift columns. (For example, a tab at column 0 occupies 8 columns, but a tab at column 6 occupies only 2 columns.) You can avoid this problem by using spaces instead of tabs with the rectangle commands.

The buffer-specific variable indent-with-tabs controls whether Epsilon does indenting with tabs or only with spaces. Set it to 0 to make Epsilon always use spaces. This variable affects only future indenting you may do; it doesn't change your file. To replace the tabs in your file, use the untabify-buffer command.

Note that the bindings shown below for kill-rectangle, copy-rectangle, and delete-rectangle only apply when there's a highlighted rectangle.

Summary:	Ctrl-x #	mark-rectangle
	Ctrl-w	kill-rectangle
	Alt-w	copy-rectangle
	Ctrl-Alt-w	delete-rectangle
		mark-line-region
		mark-inclusive-region

4.3.5 Capitalization

Epsilon has commands that allow you to change the case of words. Each travels forward, looking for the end of a word, and changes the case of the letters it travels past. Thus, if you give these commands while inside a word, only the rest of the word potentially changes case.

The Alt-l key, lowercase-word, turns all the characters it passes to lower case. The Alt-u key, uppercase-word, turns them all to upper case. The Alt-c key, capitalize-word, capitalizes a word by making

the first letter it travels past upper case, and all the rest lower case. All these commands position point after the word operated upon.

For example, the Alt-l command would turn “wOrd” into “word”. The Alt-u command would turn it into “WORD”, and the Alt-c command would turn it into “Word”.

These commands operate on the highlighted region, if there is one. If there is no highlighted region, the commands operate on the next word and move past it, as described above. The commands work on both conventional and rectangular regions.

Summary:	Alt-c	capitalize-word
	Alt-l	lowercase-word
	Alt-u	uppercase-word

4.3.6 Replacing

The key Alt-& runs the command `replace-string`, and allows you to change all occurrences of a string in the rest of your document to another string. Epsilon prompts for the string to replace, and what to replace it with. Terminate the strings with `<Enter>`. After you enter both strings, Epsilon replaces all occurrences of the first string after point with instances of the second string (but respecting any narrowing restriction; see page 176).

When entering the string to search for, you can use any of the searching subcommands described on page 45: Ctrl-c toggles case-folding, Ctrl-w toggles word searching, and Ctrl-t toggles interpreting the string as a regular expression.

To enter special characters in either the search or replace strings, use Ctrl-q before each. Type Ctrl-q Ctrl-c to include a Ctrl-c character. Type Ctrl-q Ctrl-j to include a `<Newline>` character in a search string or replacement text. Press Alt-g when entering the replacement string to copy the search string.

The key Alt-r runs the command `query-replace`, which works like `replace-string`. Instead of replacing everything automatically, however, the command positions point after each occurrence of the old string and waits for you to press a key. You may choose whether to replace this occurrence or not:

y or Y or `<Space>` Replace it, go on to next occurrence.

n or N or `<Backspace>` Don't replace it, go on to next occurrence.

! Replace all remaining occurrences. The `replace-string` command works like the `query-replace` command followed by pressing `'!'` when it shows you the first match.

`<Esc>` Exit and leave point at the match in the buffer.

`^` Back up to the previous match.

`<Period>` Replace this occurrence and then exit.

`<Comma>` Replace and wait for another command option without going on to the next match.

Ctrl-r Enter a recursive edit. Point and mark go around the match. You may edit arbitrarily. When you exit the recursive edit with Ctrl-x Ctrl-z, Epsilon restores the old mark, and the `query-replace` continues from the current location.

Ctrl-g Exit and restore point to its original location.

Ctrl-t Toggle regular expression searching. See the next section for an explanation of regular expressions.

Ctrl-w Toggle word searching.

Ctrl-c Toggle case folding.

? or help key Provide help, including a list of these options.

anything else Exit the replacement, staying at the current location, and execute this key as a command.

The command `regex-replace` operates like `query-replace`, but starts up in regular expression mode. See page 74.

The command `reverse-replace` operates like `query-replace`, but moves backwards. You can also trigger a reverse replacement by pressing `Ctrl-r` while entering the search text for any of the replacing commands.

If you invoke any of the replacing commands above with a numeric argument, Epsilon will use word searching.

If you highlight a region before replacing, Epsilon uses it as an initial search string if it's not very long. Set the `replace-in-region` variable to make Epsilon instead restrict its replacements to the highlighted region. Also see the `search-defaults-from` variable.

Replace commands preserve case. Epsilon examines the case of each match. If a match is entirely upper case, or all words are capitalized, Epsilon makes the replacement text entirely upper case or capitalized, as appropriate. Epsilon only does this when searching is case-insensitive, and neither the search string nor the replace string contain upper case letters. For example, if you search for the regular expression `welcome|hello` and replace it with `greetings`, Epsilon replaces `HELLO` with `GREETINGS` and `Welcome` with `Greetings`. See the `replace-by-case` variable to alter the rules Epsilon uses. With a regular expression replace, you can force parts of the replacement to a particular case; see page 75.

The `file-query-replace` command on `Shift-F7` replaces text in multiple files. It prompts for the search text, replacement text, and a file name which may contain wildcards. You can use extended file patterns to replace in files from multiple directories; see page 138. Epsilon skips over any file with an extension listed in `grep-ignore-file-extensions` or meeting other criteria, just like the `grep` command. See page 48 for details. To search without replacing, see the `grep` command on page 48.

With a numeric argument, this command searches through buffers instead of files. Instead of prompting for a file name pattern, Epsilon prompts for a buffer name pattern, and only operates on those buffers whose names match that pattern. Buffer name patterns use a simplified file name pattern syntax: `*` matches zero or more characters, `?` matches any single character, and character classes like `[a-z]` may be used too.

The command `delete-matching-lines` prompts for a regular expression pattern. It then deletes all lines after point in the current buffer that contain the pattern. The similar command `keep-matching-lines` deletes all lines *except* those that contain the pattern. As with any searching command, you can press `Ctrl-t`, `Ctrl-w`, or `Ctrl-c` while typing the pattern to toggle regular expression mode, word mode, or case folding (respectively).

When you select a replacing command from the menu or tool bar (rather than via a command's keyboard binding), Epsilon for Windows runs the `dialog-replace` or `dialog-regex-replace` command, to display a replace dialog. Controls on the dialog replace many of the keys described above.

Summary:	<code>Alt-&</code>	<code>replace-string</code>
	<code>Alt-r</code> , <code>Alt-%</code>	<code>query-replace</code>
	<code>Shift-F7</code>	<code>file-query-replace</code>
	<code>Alt-*</code>	<code>regex-replace</code>
		<code>reverse-replace</code>
		<code>delete-matching-lines</code>
		<code>keep-matching-lines</code>

4.3.7 Regular Expressions

Most of Epsilon's searching commands, described on page 45, take a simple string to search for. Epsilon provides a more powerful regular expression search facility, and a regular expression replace facility.

Instead of a simple search string, you provide a pattern, which describes a set of strings. Epsilon searches the buffer for an occurrence of one of the strings contained in the set. You can think of the pattern as generating a (possibly infinite) set of strings, and the regex search commands as looking in the buffer for the first occurrence of one of those strings.

The following characters have special meaning in a regex search: vertical bar, parentheses, plus, star, question mark, square brackets, period, dollar, percent sign, left angle bracket ('<'), and caret ('^'). To match them literally, they must be quoted; see page 67. See the following sections for syntax details and additional examples.

<code>abc def</code>	Finds either <code>abc</code> or <code>def</code> .
<code>(abc)</code>	Finds <code>abc</code> .
<code>abc+</code>	Finds <code>abc</code> or <code>abcc</code> or <code>abccc</code> or . . .
<code>abc*</code>	Finds <code>ab</code> or <code>abc</code> or <code>abcc</code> or <code>abccc</code> or . . .
<code>abc?</code>	Finds <code>ab</code> or <code>abc</code> .
<code>[abcx-z]</code>	Finds any single character of <code>a</code> , <code>b</code> , <code>c</code> , <code>x</code> , <code>y</code> , or <code>z</code> .
<code>[^abcx-z]</code>	Finds any single character except <code>a</code> , <code>b</code> , <code>c</code> , <code>x</code> , <code>y</code> , or <code>z</code> .
<code>.</code>	Finds any single character except (Newline).
<code>abc\$</code>	Finds <code>abc</code> that occurs at the end of a line.
<code>^abc</code>	Finds <code>abc</code> that occurs at the beginning of a line.
<code>%^abc</code>	Finds a literal <code>^abc</code> .
<code><Tab></code>	Finds a (Tab) character.
<code><#123></code>	Finds the character with ASCII code 123.
<code><p:cyrillic></code>	Finds any character with that Unicode property.
<code><alpha 1-5&!x-z></code>	Finds any alpha character except <code>x</code> , <code>y</code> or <code>z</code> or digit 1–5.
<code><^c:*comment>printf</code>	Finds uses of <code>printf</code> that aren't commented out.
<code><h:0d 0a 45></code>	Finds char sequence with those hexadecimal codes.

Figure 4.2: Summary of regular expression characters.

PLAIN PATTERNS.

In a regular expression, a string that does not contain any of the above characters denotes the set that contains precisely that one string. For example, the regular expression `abc` denotes the set that contains, as its only member, the string `'abc'`. If you search for this regular expression, Epsilon will search for the string `'abc'`, just as in a normal search.

ALTERNATION.

To include more than one string in the set, you can use the vertical bar character. For example, the regular expression `abc|xyz` denotes the set that contains the strings `'abc'` and `'xyz'`. If you search for that pattern, Epsilon will find the first occurrence of either `'abc'` or `'xyz'`. The alternation operator (`|`) always applies as widely as possible, limited only by grouping parentheses.

GROUPING.

You can enclose any regular expression in parentheses, and the resulting expression refers to the same

set. So searching for `(abc|xyz)` has the same effect as searching for `abc|xyz`, which works as in the previous paragraph. You would use parentheses for grouping purposes in conjunction with some of the operators described below.

Parentheses are also used for retrieving specific portions of the match. A regular expression replacement uses the syntax `#3` to refer to the third parenthesized group, for instance. The `find_group()` function provides a similar function for EEL programmers. The special syntax `(?:)` provides grouping just like `()`, but isn't counted as a group when retrieving parts of the match in these ways.

CONCATENATION.

You can concatenate two regular expressions to form a new regular expression. Suppose the regular expressions `p` and `q` denote sets `P` and `Q`, respectively. Then the regular expression `pq` denotes the set of strings that you can make by concatenating, to members of `P`, strings from the set `Q`. For example, suppose you concatenate the regular expressions `(abc|xyz)` and `(def|ghi)` to yield `(abc|xyz)(def|ghi)`. From the previous paragraph, we know that `(abc|xyz)` denotes the set that contains 'abc' and 'xyz'; the expression `(def|ghi)` denotes the set that contains 'def' and 'ghi'. Applying the rule, we see that `(abc|xyz)(def|ghi)` denotes the set that contains the following four strings: 'abcdef', 'abcghi', 'xyzdef', 'xyzghi'.

CLOSURE.

Clearly, any regular expression must have finite length; otherwise you couldn't type it in. But because of the closure operators, the set to which the regular expression refers may contain an infinite number of strings. If you append plus to a parenthesized regular expression, the resulting expression denotes the set of one or more repetitions of that string. For example, the regular expression `(ab)+` refers to the set that contains 'ab', 'abab', 'ababab', 'abababab', and so on. Star works similarly, except it denotes the set of zero or more repetitions of the indicated string.

OPTIONALITY.

You can specify the question operator in the same place you might put a star or a plus. If you append a question mark to a parenthesized regular expression, the resulting expression denotes the set that contains that string, and the empty string. You would typically use the question operator to specify an optional subpart of the search string.

You can also use the plus, star, and question-mark operators with subexpressions, and with non-parenthesized things. These operators always apply to the smallest possible substring to their left. For example, the regular expression `abc+` refers to the set that contains 'abc', 'abcc', 'abccc', 'abcccc', and so on. The expression `a(bc)*d` refers to the set that contains 'ad', 'abcd', 'abcbcd', 'abcbcbcd', and so on. The expression `a(b?c)*d` denotes the set that contains all strings that start with 'a' and end with 'd', with the inside consisting of any number of the letter 'c', each optionally preceded by 'b'. The set includes such strings as 'ad', 'acd', 'abcd', 'abccccbcd'.

Entering Special Characters

In a regular expression, the percent (%) character quotes the next character, removing any special meaning that character may have. For example, the expression `x%+` refers to the string 'x+', whereas the pattern `x+` refers to the set that contains 'x', 'xx', 'xxx', and so on.

You can also quote characters by enclosing them in angle brackets. The expression `x<+>` refers to the string 'x+', the same as `x%+`. In place of the character itself, you can provide the name of the character inside the angle brackets. Figure 4.3 lists some of the character names Epsilon recognizes; you can also use any character name in the Unicode standard, such as `<Superscript two>`.

<Comma>	,	<Nul>	^@	<Period>	.
<Space>		<Star>	*	<Plus>	+
<Enter>	^M	<Percent>	%	<Vbar>	
<Return>	^M	<Lparen>	(<Question>	?
<Newline>	^J	<Rparen>)	<Query>	?
<Linefeed>	^J	<Langle>	<	<Caret>	^
<Tab>	^I	<Rangle>	>	<Dollar>	\$
<Bell>	^G	<LSquare>	[<Bang>	!
<Backspace>	^H	<RSquare>]	<Exclamation>	!
<FormFeed>	^L	<Lbracket>	[<Quote>	'
<Esc>	^[<Rbracket>]	<SQuote>	'
<Escape>	^[<Dot>	.	<DQuote>	"
<Null>	^@	<Backslash>	\	<Tilde>	~

Figure 4.3: Character mnemonics in regular expressions.

To search for the NUL character (the character with ASCII code 0), use the expression <Nul>, because an actual NUL character may not appear in a regular expression.

Instead of the character's name, you can provide its numeric value using the notation <#number>. The sequence <#number> denotes the character with ASCII code *number*. For example, the pattern <#0> provides another way to specify the NUL character, and the pattern `abc<#10>+` specifies the set of strings that begin with 'abc' and end with one or more newline characters (newline has ASCII value 10). You can enter the value in hexadecimal, octal, or binary by prefixing the number with '0x', '0o', or '0b', respectively. For example, <#32>, <#0x20>, <#0o40>, and <#0b100000> all yield a <Space> character (ASCII code 32).

Character Classes

In place of any letter, you can specify a *character class*. A character class consists of a sequence of characters between square brackets. For example, the character class `[adef]` stands for any of the following characters: 'a', 'd', 'e', or 'f'.

In place of a letter in a character class, you can specify a range of characters using a hyphen: the character class `[a-m]` stands for the characters 'a' through 'm', inclusively. The class `[ae-gr]` stands for the characters 'a', 'e', 'f', 'g', or 'r'. The class `[a-zA-Z0-9]` stands for any alphanumeric character.

To specify the complement of a character class, put a caret as the first character in the class. Using the above examples, the class `[^a-m]` stands for any character other than 'a' through 'm', and the class `[^a-zA-Z0-9]` stands for any non-alphanumeric character. Inside a character class, only ^ (when it's the first character) and - have special meaning. All other characters stand for themselves, including plus, star, question mark, etc.

If you need to put a right square bracket character in a character class, put it immediately after the opening left square bracket, or in the case of an inverted character class, immediately after the caret. For example, the class `[]x]` stands for the characters ']' or 'x', and the class `[^]x]` stands for any character other than ']' or 'x'.

To include the hyphen character - in a character class, it must be the first character in the class, except for ^ and]. For example, the pattern `[^-q]` matches any character except], -, or q.

Any regular expression you can write with character classes you can also write without character classes. But character classes sometimes let you write much shorter regular expressions.

The period character (outside a character class) represents any character except a `<Newline>`. For example, the pattern `a.c` matches any three-character sequence on a single line where the first character is ‘a’ and the last is ‘c’.

You can also specify a character class using a variant of the angle bracket syntax described in the previous section for entering special characters. The expression `<Comma|Period|Question>` represents any one of those three punctuation characters. The expression `<a-z|A-Z|?>` represents either a letter or a question mark, the same as `[a-zA-Z]|<?>`, for example. The expression `<^Newline>` represents any character except newline, just as the period character by itself does.

You can also use a few character class names that match some common sets of characters.

Class	Meaning
<code><digit></code>	A digit, 0 to 9.
<code><alpha></code>	A letter, according to <code>isalpha()</code> .
<code><alphanum></code>	Either of the above.
<code><word></code>	All of the above, plus the <code>_</code> character.
<code><ident></code>	Same as <code><word></code> .
<code><hspace></code>	<code><Space Tab></code> plus Unicode spaces.
<code><wspace></code>	<code><hspace></code> plus <code><Newline></code> .
<code><ascii></code>	An ASCII character, one with a code below 128.
<code><any></code>	Any character including <code><Newline></code> .

Figure 4.4: Character Class Names

The character class `<hspace>` includes `<Space>`, and `<Tab>` plus all Unicode characters in category Z (separators). The name `<wspace>` includes all those plus `<Newline>`. Similarly, the character classes for digits, letters, and so forth include all Unicode characters of the appropriate category.

You can match all characters with a particular Unicode property, using the syntax `<p:hex-digit>`. After the `p:` part, you can put the name of a binary property as in `p:ASCIIScript`, a script name as in `p:Cyrillic`, or a category name as in `p:Zs` or `p:L`. Or you can put the name of an enumerated property, an equal sign, and a value for that property, like `p:block=Dingbats` or `p:Line_break=Alphabetic`. Case isn’t significant in these names, and certain characters like hyphen and underscore are ignored in property names.

You can combine character classes using addition, subtraction, or intersection. Addition means a matching character can be in either of two classes, as in `<alpha|digit>` to match either alphabetic characters or digits. Intersection means a matching character must be a member of both classes, as in `<p:HexDigit&p:numeric-type=decimal>`, which matches characters with the `HexDigit` binary Unicode property that also have a `Numeric-Type` property of `Decimal`. Subtraction means a matching character must be a member of one class but not another, as in `<p:currency-symbol&!dollar sign&!cent sign>` which matches all characters with the `Currency-Symbol` property except for the dollar sign and cent sign characters.

More precisely, we can say that inside the angle brackets you can put one or more character “rules”, each separated from the next by either a vertical bar `|` to add the rules together or `&` to intersect the rules. Any rule may have a `!` before it to invert that one rule, or you can put a `^` just after the opening `<` to invert the entire expression and match its complement.

Each character rule may be a character specification or a range, a character class name from the table above, or a Unicode property specification using the `p:` syntax above. A range means two character specifications with a hyphen between them. And a character specification means either the name of a

character, or # and the numeric code for a character, or the character itself (for any character except >, |, -, or \Nul).

Separately, Epsilon recognizes the syntax <h:0d 0a 45> as a shorthand to search for a series of characters by their hexadecimal codes. This example is equivalent to the pattern <#0x0d><#0x0a><#0x45>.

Regular Expression Examples

- The pattern `if|else|for|do|while|switch` specifies the set of statement keywords in C and EEL.
- The pattern `c[ad]+r` specifies strings like ‘car’, ‘cdr’, ‘caadr’, ‘caaadar’. These correspond to compositions of the car and cdr Lisp operations.
- The pattern `c[ad][ad]?[ad]?[ad]?r` specifies the strings that represent up to four compositions of car and cdr in Lisp.
- The pattern `[a-zA-Z]+` specifies the set of all sequences of 1 or more letters. The character class part denotes any upper- or lower-case letter, and the plus operator specifies one or more of those.
Epsilon’s commands to move by words accomplish their task by performing a regular expression search. They use a pattern similar to `[a-zA-Z0-9_]+`, which specifies one or more letters, digits, or underscore characters. (The actual pattern includes national characters as well.)
- The pattern `(<Newline>|<Return>|<Tab>|<Space>)+` specifies nonempty sequences of the whitespace characters newline, return, tab, and space. You could also write this pattern as `<Newline|Return|Tab|Space>+` or as `<Wspace|Return>+`, using a character class name.
- The pattern `/*.*.*/` specifies a set that includes all 1-line C-language comments. The percent character quotes the first and third stars, so they refer to the star character itself. The middle star applies to the period, denoting zero or more occurrences of any character other than newline. Taken together then, the pattern denotes the set of strings that begin with “slash star”, followed by any number of non-newline characters, followed by “star slash”. You can also write this pattern as `/<Star>.*<Star>/`.
- The pattern `/**(.|<Newline>)**/` looks like the previous pattern, except that instead of ‘.’, we have `(.|<Newline>)`. So instead of “any character except newline”, we have “any character except newline, or newline”, or more simply, “any character at all”. This set includes all C comments, with or without newlines in them. You could also write this as `/**<Any>*/*` instead.
- The pattern `<^digit|a-f>` matches any character except one of these: 0123456789abcdef.
- The pattern `<alpha&!r&!x-z&!p:softdotted>` matches all Latin letters except R, X, Y, Z, I and J (the latter two because the Unicode property SoftDotted, indicating a character with a dot that can be replaced by an accent, matches I and J). It also matches all non-Latin Unicode letters that don’t have this property.

AN ADVANCED EXAMPLE.

Let’s build a regular expression that includes precisely the set of legal strings in the C programming language. All C strings begin and end with double quote characters. The inside of the string denotes a sequence of characters. Most characters stand for themselves, but newline, double quote, and backslash must appear after a “quoting” backslash. Any other character may appear after a backslash as well.

We want to construct a pattern that generates the set of all possible C strings. To capture the idea that the pattern must begin and end with a double quote, we begin by writing

```
"something"
```

We still have to write the *something* part, to generate the inside of the C strings. We said that the inside of a C string consists of a sequence of characters. The star operator means “zero or more of something”. That looks promising, so we write

```
"(something)*"
```

Now we need to come up with a *something* part that stands for an individual character in a C string. Recall that characters other than newline, double quote, and backslash stand for themselves. The pattern `<^Newline|\"|\>` captures precisely those characters. In a C string, a “quoting” backslash must precede the special characters (newline, double quote, and backslash). In fact, a backslash may precede any character in a C string. The pattern `\(.|<Newline>)` means, precisely “backslash followed by any character”. Putting those together with the alternation operator (`|`), we get the pattern `<^Newline|\"|\>|\\(.|<Newline>)` which generates either a single “normal” character or any character preceded by a backslash. Substituting this pattern for the *something* yields

```
"(<^Newline|\"|\>|\\(.|<Newline>))*"
```

which represents precisely the set of legal C strings. In fact, if you type this pattern into a regex-search command (described below), Epsilon will find the next C string in the buffer.

Searching Rules

Thus far, we have described regular expressions in terms of the abstract set of strings they generate. In this section, we discuss how Epsilon uses this abstract set when it does a regular expression search.

When you tell Epsilon to perform a forward regex search, it looks forward through the buffer for the first occurrence in the buffer of a string contained in the generated set. If no such string exists in the buffer, the search fails.

There may exist several strings in the buffer that match a string in the generated set. Which one qualifies as the first one? By default, Epsilon picks the string in the buffer that begins before any of the others. If there exist two or more matches in the buffer that begin at the same place, Epsilon by default picks the longest one. We call this a first-beginning, longest match. For example, suppose you position point at the beginning of the following line,

```
When to the sessions of sweet silent thought
```

then do a regex search for the pattern `s[a-z]*`. That pattern describes the set of strings that start with ‘s’, followed by zero or more letters. We can find quite a few strings on this line that match that description. Among them:

```
When to the sessions of sweet silent thought
When to the sessions of sweet silent thought
When to the sessions of sweet silent thought
When to the sessions of sweet silent thought
When to the sessions of sweet silent thought
When to the sessions of sweet silent thought
```

Here, the underlined sections indicate portions of the buffer that match the description “s followed by a sequence of letters”. We could identify 31 different occurrences of such strings on this line. Epsilon picks a match that begins first, and among those, a match that has maximum length. In our example, then, Epsilon would pick the following match:

When to the sessions of sweet silent thought

since it begins as soon as possible, and goes on for as long as possible. The search would position point after the final ‘s’ in ‘sessions’.

In addition to the default first-beginning, longest match searching, Epsilon provides three other regex search modes. You can specify first-beginning or first-ending searches. For each of these, you can specify shortest or longest match matches. Suppose, with point positioned at the beginning of the following line

I summon up remembrance of things past,

you did a regex search with the pattern `m.*c|I.*t`. Depending on which regex mode you chose, you would get one of the four following matches:

I summon up remembrance of things past,	(first-ending shortest)
I <u>summon up remembrance</u> of things past,	(first-ending longest)
<u>I summon up remembrance</u> of things past,	(first-beginning shortest)
<u>I summon up remembrance of things past,</u>	(first-beginning longest)

By default, Epsilon uses first-beginning, longest matching. You can include directives in the pattern itself to tell Epsilon to use one of the other techniques. If you include the directive `<Min>` anywhere in the pattern, Epsilon will use shortest-matching instead of longest-matching. Putting `<FirstEnd>` selects first-ending instead of first-beginning. You can also put `<Max>` for longest-matching, and `<FirstBegin>` for first-beginning. These last two might come in handy if you’ve changed Epsilon’s default regex mode. The sequences `<FE>` and `<FB>` provide shorthand equivalents for `<FirstEnd>` and `<FirstBegin>`, respectively. As an example, you could use the following patterns to select each of the matches listed in the previous example:

<code><FE><Min>m.*c I.*t</code>		(first-ending shortest)
<code><FE><Max>m.*c I.*t</code>	or	<code><FE>m.*c I.*t</code> (first-ending longest)
<code><FB><Min>m.*c I.*t</code>	or	<code><Min>m.*c I.*t</code> (first-beginning shortest)
<code><FB><Max>m.*c I.*t</code>	or	<code>m.*c I.*t</code> (first-beginning longest)

You can change Epsilon’s default regex searching mode. To make Epsilon use, by default, first-ending searches, set the variable `regex-shortest` to a nonzero value. To specify first-ending searches, set the variable `regex-first-end` to a nonzero value. (Examples of regular expression searching in this documentation assume the default settings.)

When Epsilon finds a regex match, it sets point to the end of the match. It also sets the variables `matchstart` and `matchend` to the beginning and end, respectively, of the match. You can change what Epsilon considers the end of the match using the ‘!’ directive. For example, if you searched for ‘I s!ought’ in the following line, Epsilon would match the underlined section:

I sigh the lack of many a thing I sought,

Without the ‘!’ directive, the match would consist of the letters “I sought”, but because of the ‘!’ directive, the match consists of only the indicated section of the line. Notice that the first three characters of the line

also consist of ‘I s’, but Epsilon does not count that as a match. There must first exist a complete match in the buffer. If so, Epsilon will then set point and matchend according to any ‘!’ directive.

OVERGENERATING REGEX SETS.

You can use Epsilon’s regex search modes to simplify patterns that you write. You can sometimes write a pattern that includes more strings than you really want, and rely on a regex search mode to cut out strings that you don’t want.

For example, recall the earlier example of `/**(. |<Newline>)**/`. This pattern generates the set of all strings that begin with `/*` and end with `*/`. This set includes all the C-language comments, but it includes some additional strings as well. It includes, for example, the following illegal C comment:

```
/* inside /* still inside */ outside */
```

In C, a comment begins with `/*` and ends with the *very next* occurrence of `*/`. You can effectively get that by modifying the above pattern to specify a first-ending, longest match, with `<FE><Max>/**(. |<Newline>)**/`. It would match:

```
/* inside /* still inside */ outside */
```

In this example, you could have written a more complicated regular expression that generated precisely the set of legal C comments, but this pattern proves easier to write.

Regular Expression Assertions

You can force Epsilon to reject any potential match that does not line up appropriately with a line boundary, by using the ‘^’ and ‘\$’ assertions. A ‘^’ assertion specifies a beginning-of-line match, and a ‘\$’ assertion specifies an end-of-line match. For example, if you search for `^new|waste` in the following line, it would match the indicated section:

```
And with old woes new wail my dear times’s waste;
```

Even though the word ‘new’ occurs before ‘waste’, it does not appear at the beginning of the line, so Epsilon rejects it.

Other assertions use Epsilon’s angle-bracket syntax. Like the assertions ^ and \$, these don’t match any specific characters, but a potential match will be rejected if the assertion isn’t true at that point in the pattern.

Assertion	Meaning
^	At the start of a line.
\$	At the end of a line.
<bob> or <bof>	At the start of the buffer.
<eob> or <eof>	At the end of the buffer.

For example, searching for `<bob>sometext<eob>` won’t succeed unless the buffer contains only the eight character string `sometext`.

You can create new assertions from character classes specified with the angle bracket syntax by adding `[,]` or `/` at the start of the pattern.

Assertion	Meaning
<[<i>class</i> >	The next character matches <i>class</i> , the previous one does not.
<][<i>class</i> >	The previous character matches <i>class</i> , the next one does not.
</ <i>class</i> >	Either of the above.

The *class* in the above syntax is a |-separated or &-separated list of one or more single characters, character names like Space or Tab, character numbers like #32 or #9, ranges of any of these, character class names like Word or Digit, or Unicode property specifications. See page 69 for details on character classes.

For example, `</word>` matches at a word boundary, and `<]word>` matches at the end of a word. The pattern `<]0-9|a-f>` matches at the end of a run of hexadecimal digits. The pattern `(cat|[0-9])</digit>(dog|[0-9])` matches `cat3` or `4dog`, but not `catdog` or `42`. The pattern `<[p:cyrillic>` matches at the start of a run of Cyrillic characters.

COLOR CLASS ASSERTIONS.

Another type of assertion matches based on the next character's color class for syntax highlighting. `<^c:*comment>printf` finds uses of `printf` that aren't commented out. `<[c:perl-string>` finds "characters that start a string in Perl mode, ignoring those that end it, or appear quoted inside it, or in comments or other places.

The text after the `c:` is a simple filename-style pattern that will be matched against the name of the color class: `*` matches zero or more characters, `?` matches any single character, and simple ranges with `[]` are allowed. A character with no syntax highlighting applied will match the name "none". This type of assertion may start with `^` to invert the matching rules, or with `/`, `[` or `]` to match color boundaries.

To apply more than one assertion to a character, put them in sequence.

`<^c:perl-string><^c:*comment>printf` finds instances of `printf` that are in neither Perl strings nor comments.

You can use the `set-color` command to see the color class names Epsilon uses.

When you combine assertions with operators `*` or `+`, you must use parentheses to specify that the assertion applies to each character. `(<^c:*-comment><any>)+` matches a run of non-comment characters. Without the parentheses the assertion only applies to the first character of the run.

In extension language code, use the `do_color_searching()` subroutine if your regular expression might include syntax highlighting assertions, which ensures the buffer's syntax highlighting is up to date.

Regular Expression Commands

You can invoke a forward regex search with the Ctrl-Alt-s key, which runs the command `regex-search`. The Ctrl-Alt-r key invokes a reverse incremental search. You can also enter regular expression mode from any search prompt by typing Ctrl-t to that prompt. For example, if you press Ctrl-s to invoke `incremental-search`, pressing Ctrl-t causes it to enter regular expression mode. See page 45 for a description of the searching commands.

The key Alt-* runs the command `regex-replace`. This command works like the command `query-replace`, but interprets its search string as a regular expression.

In the replacement text of a regex replace, the `#` character followed by a digit *n* from 1 to 9 has a special meaning in the replacement text. Epsilon finds the *n*th parenthesized expression in the pattern, counting left parentheses from 1. It then substitutes the match of this subpattern for the `#n` in the replacement text. For example, replacing

```
([a-zA-Z0-9_]+) = ([a-zA-Z0-9_]+)
```

with

```
#2 := #1
```

changes

```
variable = value;
```

to

```
value := variable;
```

If #0 appears in the replacement text, Epsilon substitutes the entire match for the search string. To include the actual character # in a replacement text, use ##. In a search pattern, you can follow the open parenthesis with ?: to tell Epsilon not to count it for replacement purposes; that pair of parentheses will only be used for grouping.

The replacement text can use the syntax #U to force the rest of the replacement to uppercase (including text substituted from the match using #1 syntax). Using #L or #C forces the remaining text to lowercase, or capitalizes it, respectively. Using #E marks the end of such case modifications; the following replacement text will be substituted as-is. For instance, searching for “(<word>+) by (<word>+)” and replacing it with “#L#2#E By #U#1” will change the match “Two by Four” into “four By TWO”.

When the search string consists of multiple words of literal text separated by the | character, you can use #S in the replacement text to swap them. For instance, if you search for dog|cat and replace it with #S, Epsilon replaces instances of dog with cat, and instances of cat with dog. If you have more than two choices, each choice will be replaced by the next choice in the list.

When you don’t use the above syntax, replacing preserves the case of each match according to specific rules. See the `replace-by-case` variable for details.

Characters other than # in the replacement text have no special meaning. To enter special characters, type a Ctrl-q before each. Type Ctrl-q Ctrl-j to include a `<Newline>` character in the replacement text. Or specific characters in the replacement text by name, using the syntax #<Newline>, or by number, such as #<#0x221a> for the Unicode square root character.

Summary:	Ctrl-Alt-s	regex-search
	Ctrl-Alt-r	reverse-regex-search
	Alt-*	regex-replace

4.3.8 Rearranging

Sorting

Epsilon provides several commands to sort buffers, or parts of buffers.

The `sort-buffer` command lets you sort the lines of the current buffer. The command asks for the name of a buffer in which to place the sorted output. The `sort-region` command sorts the part of the current buffer between point and mark, in place. The commands `reverse-sort-buffer` and `reverse-sort-region` operate like the above commands, but reverse the sorting order.

By default, all the sorting commands sort the lines by considering all the characters in the line. If you prefix a numeric argument of *n* to any of these commands, they will compare lines starting at column *n*.

When comparing lines of text during sorting, Epsilon normally folds lower case letters to upper case before comparison, if the `case-fold` variable has a nonzero value. If the `case-fold` variable has a value of 0, Epsilon compares characters as-is. However, setting the buffer-specific `sort-case-fold` variable to 0 or 1 overrides the `case-fold` variable, for sorting purposes. By default, `sort-case-fold` has a value of 2, which means to defer to `case-fold`.

Summary:	sort-buffer
	sort-region

reverse-sort-buffer
reverse-sort-region

Comparing By Lines

Like `compare-windows` and `diff`, the `compare-sorted-windows` command compares the contents of the current buffer with that displayed in the next window on the screen, but it generates lists of lines that appear in both buffers or just one.

Use it when you have (for example) two lists of variable names, and you want to find out which variables appear on only one or the other list, and which appear on both. This command assumes that you sorted both the buffers. It copies all lines appearing in both buffers to a buffer named “inboth”. It copies all lines that appear only in the first buffer to a buffer named “only1”, and lines that appear only in the second to a buffer named “only2”.

The `uniq` command goes through the current buffer and looks for adjacent identical lines, deleting the duplicate copies of each repeated line and leaving just one. It doesn’t modify any lines that only occur once. This command behaves the same as the Unix command of the same name.

The `keep-unique-lines` command deletes all copies of any duplicated lines. This command acts like the Unix command “`uniq -u`”.

The `keep-duplicate-lines` command deletes all lines that only occur once, and leaves one copy of each duplicated line. This command acts like the Unix command “`uniq -d`”.

The following table shows how sample text would be modified by each of the above commands.

Sample text	Uniq	Keep-duplicate-lines	Keep-unique-lines
dog	dog	dog	cat
dog	cat	horse	rabbit
cat	horse		dog
horse	rabbit		
horse	dog		
horse			
rabbit			
dog			

Summary:

`compare-sorted-windows`
`uniq`
`keep-unique-lines`
`keep-duplicate-lines`

Transposing

Epsilon has commands to transpose characters, words, and lines. To transpose the words before and after point, use the `Alt-t` command. This command leaves undisturbed any non-word characters between the words. Point moves between the words. The `Ctrl-x Ctrl-t` command transposes the current and previous lines and moves point between them.

The Ctrl-t command normally transposes the characters before and after point. However, at the start of a line it transposes the first two characters on the line, and at the end of a line it transposes the last two. On a line with one or no characters, it does nothing.

Summary:	Ctrl-t	transpose-characters
	Alt-t	transpose-words
	Ctrl-x Ctrl-t	transpose-lines

Formatting Text

Epsilon has some commands that make typing manuscript text easier.

You can change the right margin, or *fill column*, using the Ctrl-x f command. By default, it has a value of 70. With a numeric argument, the command sets the fill column to that column number. Otherwise, this command tells you the current value of the fill column and asks you for a new value. If you don't provide a new value but instead press the `<Enter>` key, Epsilon will use the value of point's current column. For example, you can set the fill column to column 55 by typing Ctrl-u 55 Ctrl-x f. Alternatively, you can set the fill column to point's column by typing Ctrl-x f `<Enter>`. The buffer-specific variable `margin-right` stores the value of the fill column. To set the default value for new buffers you create, use the `set-variable` command on F8 to set the default value of the `margin-right` variable. (See the `c-fill-column` variable for the C mode equivalent.) A file's contents can specify a particular fill column; see page 127.

In *auto fill mode*, you don't have to worry about typing `<Enter>`'s to go to the next line. Whenever a line gets too long, Epsilon breaks the line at the appropriate place if needed. The `auto-fill-mode` command enables or disables auto filling (word wrap) for the current buffer. With a numeric argument of zero, it turns auto filling off; with a nonzero numeric argument, it turns auto filling on. With no numeric argument, it toggles auto filling. During auto fill mode, Epsilon shows the word "Fill" in the mode line. The buffer-specific variable `fill-mode` controls filling. If it has a nonzero value, filling occurs. To make Epsilon always use auto fill mode, you can use the `set-variable` command to set the default value of `fill-mode`.

In some language modes, Epsilon uses a special version of auto-fill mode that typically only fills text in certain types of comments. See page 104 for details.

Epsilon normally indents new lines it inserts via auto fill mode so they match the previous line. The buffer-specific variable `auto-fill-indents` controls whether or not Epsilon does this. Epsilon indents these new lines only if `auto-fill-indents` has a nonzero value. Set the variable to 0 if you don't want this behavior.

During auto filling, the `normal-character` command first checks to see if the line extends past the fill column. If so, the extra words automatically move down to the next line.

The `<Enter>` key runs the command `enter-key`, which behaves like `normal-character`, but inserts a newline instead of the character that invoked it. Epsilon binds this command to the `<Enter>` key, because Epsilon uses the convention that Ctrl-j's separate lines, but the keyboard has the `<Enter>` key yield a Ctrl-m. In overwrite mode, the `<Enter>` key simply moves to the beginning of the next line.

The Alt-q command fills the current paragraph. The command fills each line by moving words between lines as necessary, so the lines but the last become as long as possible without extending past the fill column. If the screen shows a highlighted region, the command fills all paragraphs in the region. The `fill-region` command fills all paragraphs in the region between point and mark, whether or not the region is highlighted.

If you give a numeric prefix argument of five or less to the above filling commands, they unwrap lines in a paragraph, removing all line breaks. Alt-2 Alt-q is one quick way to unwrap the current paragraph. With a numeric argument greater than 5, the paragraph is filled using that value as a temporary right margin. (Note

that C mode places a different fill command on Alt-q, and it interprets an argument to mean “fill using the current column as a right margin”).

Alt-Shift-Q runs the prefix-fill-paragraph command. It fills the current paragraph while preserving any run of spaces, punctuation, and other non-alphanumeric characters that appears before each of the lines in the paragraph. Highlight a region first and it will fill all the paragraphs within in this manner. With a numeric argument, it fills the paragraph using the current column as the right margin, instead of the margin-right variable.

The fill-indented-paragraph command is similar; it fills the current paragraph as above, but tries to preserve only indentation before each line of the paragraph. It’s better than prefix-fill-paragraph when the paragraph to be filled contains punctuation characters and similar that should be filled as part of the paragraph, not considered part of the prefix.

The mail-fill-paragraph command on Ctrl-c Alt-q is similar to prefix-fill-paragraph, but specialized for the quoting rules of email that put > or # before each line. It preserves email quoting characters at the starts of lines, treating other characters as part of the paragraph.

Press Ctrl-c > to add email-style quoting to the current paragraph (or highlighted region). Press Ctrl-c < to remove such quoting.

These mail-formatting commands use the mail-quote-pattern, mail-quote-skip, and mail-quote-text variables.

Summary:	Ctrl-x f	set-fill-column
	Alt-q	fill-paragraph
	Alt-Shift-Q	prefix-fill-paragraph
	Ctrl-c Alt-q	mail-fill-paragraph
	Ctrl-c >	mail-quote-region
	Ctrl-c <	mail-unquote
		fill-indented-paragraph
		fill-region
		auto-fill-mode
	⟨Enter⟩	enter-key

4.3.9 Indenting Commands

Epsilon can help with indenting your program or other text. The ⟨Tab⟩ key runs the indent-previous command, which makes the current line start at the same column as the previous non-blank line. Specifically, if you invoke this command with point in or adjacent to a line’s indentation, indent-previous replaces that indentation with the indentation of the previous non-blank line. If point’s indentation exceeds that of the previous non-blank line, or if you invoke this command with point outside of the line’s indentation, this command simply inserts a ⟨Tab⟩. See page 110 for information on changing the width of a tab.

Epsilon can automatically indent for you when you press ⟨Enter⟩. Setting the buffer-specific variable auto-indent nonzero makes Epsilon do this. The way Epsilon indents depends on the current mode. For example, C mode knows how to indent for C programs. In Epsilon’s default mode, fundamental mode, Epsilon indents like indent-previous if you set auto-indent nonzero. (Auto-indenting removes trailing spaces and tabs too.)

In some modes Epsilon not only indents the newly inserted line, but also reindents the existing line. Variables named after their modes, like c-reindent-previous-line, control this. The default-reindent-previous-line variable controls this for modes that don’t have their own variable.

When Epsilon automatically inserts new lines for you in auto fill mode, it looks at a different variable to determine whether to indent these new lines. Epsilon indents in this case only if the buffer-specific variable `auto-fill-indents` has a nonzero value.

The `Alt-m` key moves point to the beginning of the text on the current line, just past the indentation.

The `indent-under` command functions like `indent-previous`, but each time you invoke it, it indents more, to align with the next word in the line above. In detail, it goes to the same column in the previous non-blank line, and looks to the right for the end of the next region of spaces and tabs. It indents the rest of the current line to that column after removing spaces and tabs from around point. With a highlighted region, it indents all lines in the region to that same column.

With a numeric prefix argument, `indent-under` goes to a different run of non-spaces. For instance, with an argument of 3, it goes to the previous line and finds the third word after the original column, then aligns the original line there.

The `indent-rigidly` command, bound to `Ctrl-x Ctrl-i` (or `Ctrl-x <Tab>`), changes the indentation of each line between point and mark by a fixed amount provided as a numeric argument. For instance, `Ctrl-u 8 Ctrl-x Ctrl-i` moves all the lines to the right by eight spaces. With no numeric argument, lines move to the right by the buffer's tab size (default 8; see page 110), and with a negative numeric argument, lines move to the left. So, for example, `Ctrl-u -1000 Ctrl-x Ctrl-i` should remove all the indentation from the lines between point and mark.

If you highlight a region before pressing `<Tab>` (or any key that runs one of the commands `indent-previous` or `do-c-indent`), Epsilon indents all lines in the region by one tab stop, by calling the `indent-rigidly` command. You can provide a numeric argument to specify how much indentation you want.

The `Shift-<Tab>` key moves the cursor back to the previous tab stop. But if you highlight a region before pressing it, it will remove one tab stop's worth of indentation. (See the `resize-rectangle-on-tab` variable if you want these keys to instead change the region's shape without moving text.)

The `indent-region` command, bound to `Ctrl-Alt-\`, works similarly. It goes to the start of each line between point and mark and invokes the command bound to `<Tab>`. If the resulting line then contains only spaces and tabs, Epsilon removes them.

You can set up Epsilon to automatically reindent text when you yank it. Epsilon will indent like `indent-region`. By default, Epsilon does this only for C mode (see the `reindent-after-c-yank` variable).

To determine whether to reindent yanked text, the yank command first looks for a variable whose name is derived from the buffer's mode as it appears in the mode line: `reindent-after-c-yank` for C mode buffers, `reindent-after-html-yank` for HTML mode buffers, and so forth. If there's no variable by that name, Epsilon uses the `reindent-after-yank` variable instead. Instead of a variable, you can write an EEL function with the same name; Epsilon will call it and use its return value. See the description of `reindent-after-yank` for details on what different values do.

The `Alt-s` command horizontally centers the current line between the first column and the fill column by padding the left with spaces and tabs as necessary. Before centering the line, the command removes spaces and tabs from the beginning and end of the line.

With any of these commands, Epsilon indents by inserting as many tabs as possible without going past the desired column, and then inserting spaces as necessary to reach the column. You can set the size of a tab by setting the `tab-size` variable. Set the `soft-tab-size` variable if you want Epsilon to use one setting for displaying existing tab characters, and a different one for indenting.

If you prefer, you can make Epsilon indent using only spaces. The buffer-specific variable `indent-with-tabs` controls this behavior. Set it to 0 using `set-variable` to make Epsilon use only spaces when inserting indentation.

If you want `<Tab>` to simply indent to the next tab stop, you can bind the `indent-to-tab-stop` command to it. To disable smart indenting in a particular language mode, you can bind this command to `<Tab>` only in that mode.

The `untabify-region` command on `Ctrl-x Alt-i` changes all tab characters between point and mark to the number of spaces necessary to make the buffer look the same. The `tabify-region` command on `Ctrl-x Alt-<Tab>` does the reverse. It looks at all runs of spaces and tabs, and replaces each with tabs and spaces to occupy the same number of columns. The commands `tabify-buffer` and `untabify-buffer` are similar, but operate on the entire buffer, instead of just the region.

Summary:	<code>Alt-m</code>	<code>to-indentation</code>
	<code><Tab></code>	<code>indent-previous</code>
	<code>Shift-<Tab></code>	<code>back-to-tab-stop</code>
	<code>Ctrl-Alt-i</code>	<code>indent-under</code>
	<code>Ctrl-x <Tab></code>	<code>indent-rigidly</code>
	<code>Ctrl-Alt-\</code>	<code>indent-region</code>
	<code>Alt-s</code>	<code>center-line</code>
	<code>Ctrl-x Alt-<Tab></code>	<code>tabify-region</code>
	<code>Ctrl-x Alt-i</code>	<code>untabify-region</code>
		<code>tabify-buffer</code>
		<code>untabify-buffer</code>
		<code>indent-to-tab-stop</code>

4.3.10 Aligning

The `align-region` command on `Ctrl-c Ctrl-a` aligns elements on lines within the current region. It changes the spacing just before each element, so it starts at the same column on every line where it occurs.

It uses alignment rules specialized for the current mode. By default, it aligns the first “=” character on each line, and any comments on the lines.

For C mode, Epsilon additionally aligns the names of variables being defined (in simple definitions), the definitions of macros in `#define` lines, and the backslash character at the end of preprocessor commands. It can change

```
int hour = 3;          // The hour.
short int minute = 22; // The minute.
int second = 14;       // The second.

#define GET_HOUR()  hour    // Get the hour.
#define GET_MINUTE() minute // Get the minute.
#define GET_SECOND() second // Get the second.

into

int      hour   = 3;          // The hour.
short int minute = 22;       // The minute.
int      second = 14;        // The second.

#define GET_HOUR()  hour    // Get the hour.
```

```
#define GET_MINUTE()  minute  // Get the minute.
#define GET_SECOND()  second  // Get the second.
```

You can disable individual alignment rules by setting the `align-region-rules` variable, or increase the minimum spacing used by all automatic rules by setting the `align-region-extra-space` variable.

The command can also perform alignments specified manually. Run it with a numeric prefix argument, and it will prompt for a regular expression pattern that defines the alignment rule. It must consist of two parenthesized patterns, such that in a regular expression replacement, #1 and #2 would substitute their text. Alignment will alter the spacing between these two elements. Manual alignment will also prompt for the amount of additional spacing to be added between the two elements.

To use the built-in mode-based rules, but add extra space, run `align-region` with a numeric prefix, but enter nothing for the search pattern. The command will prompt for the amount of additional space and apply it using the mode's default alignment rules, as if you had temporarily modified the `align-region-extra-space` variable.

The `align-by-tab` command modifies tab-delimited text in the region (text that uses a single tab character to mark each new column, such as text copied from a spreadsheet), reformatting it using spaces and tabs so the text appears in columns. It makes each column just wide enough for its contents.

Summary:	Ctrl-c Ctrl-a	<code>align-region</code>
		<code>align-by-tab</code>

4.3.11 Automatically Generated Text

The `copy-file-name` command on Ctrl-c Alt-n is a convenient way to put the current buffer's filename onto the clipboard. In a dired buffer, it copies the current line's absolute pathname.

The similar `copy-include-file-name` on Ctrl-c Alt-i formats the current file's name as an `#include` command for C mode buffers, and similarly for other languages. It looks for a variable with a name of the form `copy-include-file-name-mode`, where *mode* is the current mode name. The variable holds a file name template (see page 122) which is used to format the current file's name. If there's a function by that name, not a variable, Epsilon simply calls it. The function can call the `copy_line_to_clipboard()` subroutine after preparing a suitable line.

The `insert-date` command on Ctrl-c Alt-d inserts the current time and/or date, according to the format specified by the `date-format` variable.

Summary:	Ctrl-c Alt-n	<code>copy-file-name</code>
	Ctrl-c Alt-i	<code>copy-include-file-name</code>
	Ctrl-c Alt-d	<code>insert-date</code>

4.3.12 Spell Checking

The Spell minor mode makes Epsilon highlight misspelled words as you edit.

First configure spell checking by running the `spell-configure` command. The first time you run it, it will download and install a set of dictionary files into the "spell" subdirectory of your customization directory. (See <https://www.lugaru.com/spell.html> if you need to download it manually.) Then it will ask your region (American, Canadian, British, or British with -ize spellings preferred) and other questions like dictionary

size. (A larger dictionary means rarer words won't be marked as potential misspellings, but it will miss those misspellings that happen to result in rare words.) To start, just choose default options for each question.

Use the `spell-mode` command to make Epsilon highlight misspelled words in the current mode. The command toggles highlighting; a numeric prefix argument forces it on (if nonzero) or off (if zero). “Sp” in the mode line indicates Spell minor mode is on. Use the `buffer-spell-mode` command instead if you want Epsilon to only highlight misspelled words in the current buffer.

Epsilon remembers whether you want spell checking in a particular mode using a variable like `html-spell-options`, whose name is derived from the mode name. If a mode has no associated variable, Epsilon uses the `default-spell-options` variable. Each variable contains bits to further customize spelling rules for that mode. The `0x1` bit says whether misspelled words should be highlighted at all; `spell-mode` toggles it. The following table shows the meaning of the other bits in each variable.

Bit	Meaning
0x1	Highlight misspelled words.
0x2	Skip words containing an underscore.
0x4	Skip MixedCaseWords (those with internal capitalization).
0x8	Skip uppercase words (those with no lowercase letters).
0x10	Skip words following a digit, like 14th.
0x20	Skip words before a digit, like gr8.
0x200	Don't remove 's when checking words.
0x1000	Provide faster but less accurate built-in suggestions.
0x2000	Don't copy the case of the original in built-in suggestions.
0x4000	Add globally ignored words to spell helper's list.

The `spell-correct` command can suggest replacements for a misspelled word. It can also record a word in an ignore list so Epsilon no longer highlights it as a misspelling. Epsilon maintains a global ignore list named `ignore.lst` in your customizations directory. That directory also contains its main word list dictionary `espell.lst` (which is ordered so that more common words appear closer to the top of the file) and `espell.srt`, a (case-sensitively) sorted version of `espell.lst`.

Epsilon also checks directory-specific, file-specific, and mode-specific ignore lists. When checking a file named `file.html`, for example, Epsilon looks for an ignore list file named `.file.html.espell` in that same directory, and a directory-specific ignore list file in that directory named `.directory.espell`. A mode-specific ignore list file is named `ignore.modename.mode.lst`, where *modename* is the current mode name, and appears in your customization directory.

All these files contain one word per line. Epsilon automatically sorts ignore list files when it uses them. (Epsilon can optionally use extension-specific ignore lists too. By default this is disabled for simplicity. See the `global-spell-options` variable.)

The `spell-buffer-or-region` command performs spell checking for the current buffer, going to each misspelled word in turn and asking if you want to correct it or ignore it. With a highlighted region it checks just that region.

The `spell-grep` command writes a copy of all lines with spelling errors to the grep buffer, where you can use the usual grep commands to navigate among them. See page 48 for details.

MAKING SUGGESTIONS

The `spell-correct` command presents a list of suggestions. Epsilon can generate these in several different ways. The default method uses the installed dictionary files. A faster, less accurate, but still self-contained method is available by setting a bit in the current `-spell-options` variable.

Epsilon can also run an external program to provide suggestions; this is generally very fast and produces the best suggestions. The `spell-configure` command configures this. It sets up Epsilon to use `aspell` or the older `ispell`, two free command line spelling programs often installed on Unix systems. It can also set up Epsilon to use `MicroSpell`, a commercial spell checking program for Windows systems available from <http://www.microspell.com>, by installing a helper program `mspellcmd.exe` into its directory. In Epsilon for macOS, it can also use the Mac's native spelling engine, though this is not available when you run Epsilon for macOS over a network connection from another computer.

CUSTOMIZING SPELL CHECKING

Epsilon looks for words to be checked using a regular expression pattern. In modes without syntax highlighting, it uses the pattern in the `default-spell-word-pattern` variable. In modes with syntax highlighting, it uses `default-color-spell-word-pattern`.

This latter pattern makes Epsilon ignore words based on their syntax highlighting color class, so that it skips over language keywords, variable names, and so forth. It checks words only if the mode colors them using a color class whose name ends in `-text`, `-comment`, or `-string`. It uses a color class assertion (see page 74) to do this.

You can define a replacement spell check pattern for any mode by creating a variable whose name is the mode name followed by `-spell-word-pattern`. Then Epsilon will use that variable instead of one of the default variables. For instance, if you want XML mode to check attributes as well as text but not comments, you could define a variable `xml-spell-word-pattern` and copy its value from the `default-color-spell-word-pattern` variable, changing the color class assertion to `<c:*(-text|*-attributes)>`.

A mode can make the speller ignore words based on adjacent text, in addition to using color class assertions. Create a variable whose name is the mode's name followed by `-spell-ignore-pattern-prefix`. If it exists, and the regular expression pattern it contains matches the text just before a word, the speller will skip it. For instance, if in Sample mode a `#` at the start of a line indicates a comment, define a variable `sample-spell-ignore-pattern-prefix` and set it to `^#.*`. Similarly, a variable ending in `-spell-ignore-pattern-suffix` that matches just after a word will make the speller ignore the word.

A mode can define an alternative set of dictionaries and ignore files by setting the buffer-specific `spell_language_prefix` variable. Set it to a suffix like `"-fr"` and Epsilon will look for alternative files, which the mode must supply, ending with that suffix.

Summary:

	<code>spell-mode</code>
	<code>buffer-spell-mode</code>
	<code>spell-configure</code>
	<code>spell-buffer-or-region</code>
	<code>spell-grep</code>
<code>Ctrl-c Ctrl-o</code>	<code>spell-correct</code>

4.3.13 Hex Mode

The `hex-mode` command creates a second buffer that shows a hex listing of the original buffer. You can edit this buffer, as explained below. Press `q` when you're done, and Epsilon will return to the original buffer, offering to apply your changes.

A hex digit (0-9, a-f) in the left-hand column area moves in the hex listing to the new location.

A hex digit (0-9, a-f) elsewhere in the hex listing modifies the listing.

q quits hex mode, removing the hex mode buffer and returning to the original buffer. Epsilon will first offer to apply your editing changes to the original buffer.

<Tab> moves between the columns of the hex listing.

s or r searches by hex bytes. Type a series of hex bytes, like 0a 0d 65, and Epsilon will search for them. S searches forward, R in reverse.

Ctrl-s and Ctrl-r temporarily toggle to the original buffer so you can search for literal text. When the search ends, they move to the corresponding place in the hex listing.

t toggles between the original buffer and the hex mode buffer, going to the corresponding position.

prompts for a new character value and overwrites the current character with it. You can use any of these formats: 'A', 65, 0x41 (hex), 0b1100101 (binary), 0o145 (octal).

n or p move to the next or previous line.

g prompts for an offset in hexadecimal, then goes there.

o toggles the hex overwrite submode, which changes how Epsilon interprets keys you type in the rightmost column of the hex listing. In overwrite mode, printable characters you type in the rightmost column overwrite the text there, instead of acting as hex digits or commands.

For instance, typing "3as" in the last column while in overwrite mode replaces the next three characters with the characters 3, a, and s. Outside overwrite mode, they replace the current character with one whose hex code is 3a, and then begin a search.

To use hex mode commands from overwrite mode, prefix them with a Ctrl-c character, such as Ctrl-c o to exit overwrite mode. Or move out of the rightmost column with **<Tab>** or other movement keys.

? shows help on hex mode.

Summary:

hex-mode

4.4 Language Modes

When you use the `find-file` command to read in a file, Epsilon looks at the file's extension to see if it has a mode appropriate for editing that type of file. For example, when you read a `.h` file, Epsilon goes into C mode. Specifically, whenever you use `find-file` and give it a file name "`foo.ext`", after `find-file` reads in the file, it executes a command named "`suffix_ext`", if such a command exists. The `find-file` command constructs a subroutine name from the file extension to allow you to customize what happens when you begin editing a file with that extension.

For example, if you want to enter C mode automatically whenever you use `find-file` on a `.x` file, you simply create a command (a keyboard macro would do) called "`suffix_x`", and have that command call `c-mode`, or even better, an existing `suffix_` function. One way is to add a line like this to your `einit.ecm` file (see page 164):

```
(define-macro "suffix-x" "<!suffix-c>")
```

For another example, you can easily stop Epsilon from automatically entering C mode on a `.h` file by using the `delete-name` command to delete the subroutine "`suffix-h`". (You can interchange the `-` and `_`

characters in Epsilon command names.) Or define a suffix-h macro so it calls the fundamental-mode command in your `init.ecm` file, as above.

When Epsilon doesn't recognize a file's extension, a file has no extension, or when an extension is frequently used for different kinds of data, Epsilon examines the file's contents to help pick the best mode. A subroutine you can customize named `guess_mode_without_extension()` performs this task.

A file's text can include a file variable (see page 127) that specifies the correct mode to use, and Epsilon will use that mode instead of choosing one itself.

Epsilon also has various features that are useful in many different language modes. See the description of tagging on page 50 and the section starting on page 100.

In addition to the language-specific modes described in the following sections, Epsilon includes modes that support various Epsilon features. For example, the buffer listing generated by the `bufed` command on `Ctrl-x Ctrl-b` is actually in an Epsilon buffer, and that buffer is in Bufed mode. Press `F1 m` to display help on the current mode.

Many language modes will call a hook function if you've defined one. For example, C mode tries to call a function named `c_mode_hook()`. A hook function is a good place to customize a mode by setting buffer-specific variables. It can be a keyboard macro or a function written in EEL, and it will be called whenever Epsilon loads a file that should be in the specified mode.

To customize a mode's key bindings, see the example for C mode on page 88.

The fundamental-mode command removes changes to key bindings made by modes such as C mode, Dired mode, or Bufed mode. You can configure Epsilon to highlight matching parentheses and other delimiters in fundamental mode; see the `fundamental-auto-show-delim-chars` variable.

Also see page 125 to customize the list of file types shown in File/Open and similar dialogs in Epsilon for Windows.

Summary: fundamental-mode

4.4.1 Asm Mode

Epsilon automatically enters Asm mode when you read a file with an extension of `.asm`, `.inc`, `.al`, `.mac`, `.ah`, or `.asi`. In Asm mode, Epsilon does appropriate syntax highlighting, tagging, and commenting. The `compile-buffer` command uses the `compile-asm-cmd` variable in this mode.

Summary: asm-mode

4.4.2 Batch Mode

Epsilon automatically enters Batch mode when you read a file with an extension of `.bat`, `.cmd`, or `.btm`. In Batch mode, Epsilon does appropriate syntax highlighting, and provides delimiter highlighting using the `auto-show-batch-delimiters` and `batch-auto-show-delim-chars` variables.

Summary: batch-mode

4.4.3 C Mode

The `c-mode` command puts the current buffer in C mode. C mode provides smart indenting for programs written in C, C++, C#, JavaScript, Java, Epsilon's extension language EEL, Objective-C, and other C-like languages. Pressing `<Enter>` or `<Tab>` examines previous lines to find the correct indentation. Epsilon supports several common styles of indentation, controlled by some extension language variables.

The `closeback` variable controls the position of the closing brace:

<code>closeback = 0;</code>	<code>closeback = 1;</code>
<code>if (foo){</code>	<code>if (foo){</code>
<code>bar();</code>	<code>bar();</code>
<code>baz();</code>	<code>baz();</code>
<code>}</code>	<code>}</code>

By placing the opening brace on the following line, you may also use these styles:

<code>closeback = 0;</code>	<code>closeback = 1;</code>
<code>if (foo)</code>	<code>if (foo)</code>
<code>{</code>	<code>{</code>
<code>bar();</code>	<code>bar();</code>
<code>baz();</code>	<code>baz();</code>
<code>}</code>	<code>}</code>

`Closeback` by default has a value of 1. This variable also recognizes the value 2 to indicate that lines starting with a right parenthesis should be indented to match the start of the line with its matching left parenthesis, not the left parenthesis itself.

Use the `Topindent` variable to control the indentation of top-level statements in a function:

<code>Topindent = 0;</code>	<code>Topindent = 1;</code>
<code>foo()</code>	<code>foo()</code>
<code>{</code>	<code>{</code>
<code>if (bar)</code>	<code>if (bar)</code>
<code>baz();</code>	<code>baz();</code>
<code>}</code>	<code>}</code>

`Topindent` by default has a value of 1.

The `Matchdelim` variable controls whether typing `,` `]`, or `}` displays the corresponding `(`, `[`, or `{` using the `show-matching-delimiter` command. The `Matchdelim` variable normally has a value of 1, which means that Epsilon shows matching delimiters. You can change these variables as described on page 159.

In C mode, the `<Tab>` key reindents the current line if pressed with point in the current line's indentation. `<Tab>` just inserts a tab if pressed with point somewhere else, or if pressed two or more times successively. If you set the variable `c-tab-always-indents` to 1, then the `<Tab>` key will reindent the current line, regardless of your position on the line. If you press it again, it will insert another tab. The `<Enter>` key indents the line it inserts, as well as the current line (but see the `c-reindent-previous-line` variable).

When you yank text into a buffer in C mode, Epsilon automatically reindents it. This is similar to the "smart paste" feature in some other editors. You can set the variable `reindent-after-c-yank` to zero to disable this behavior. Epsilon doesn't normally reindent comments when yanking; set the

`reindent-c-comments` and `reindent-one-line-c-comments` variables to change that. Also see the `reindent-c-preprocessor-lines` variable.

By default, Epsilon uses the value of the buffer-specific `tab-size` variable to determine how far to indent. For example, if the tab size has a value of 5, Epsilon will indent the line following an `if` statement five additional columns.

If you want the width of a tab character in C mode buffers to be different than in other buffers, set the variable `c-tab-override` to the desired value. C mode will change the buffer's tab size to the specified number of columns. The `eel-tab-override` variable does the same in EEL buffers (which use a variation of C mode). Also see the description of file variables on page 127 for a way in which individual files can indicate they should use a particular tab size.

If you want to use one value for the tab size and a different one for C indentation, set the buffer-specific `c-indent` variable to the desired indentation using the `set-variable` command. When `c-indent` has a value of zero, as it has by default, Epsilon uses the `tab-size` variable for its indentation. (Actually, the `<Tab>` key in C mode doesn't necessarily insert a tab when you press it two or more times in succession. Instead, it indents according to `c-indent`. If the tab size differs from the C indent, it may have to insert spaces to reach the proper column.)

In Java files, Epsilon uses the similar variable `java-indent` to set the column width of one level of indentation.

The `c-case-offset` variable controls the indentation of case statements. Normally, Epsilon indents them one level more than their controlling `switch` statements. Epsilon adds the value of this variable to its normal indentation, though. If you normally indent by 8 spaces, for example, and want case statements to line up with their surrounding `switch` statements, set `c-case-offset` to `-8`.

Similarly, the `c-access-spec-offset` variable controls the indentation of `public:`, `private:`, `protected:` (and, for C#, `internal:`) access specifiers.

The `c-label-indent` variable provides the indentation of lines starting with labels. Normally, Epsilon moves labels to the left margin.

Epsilon offsets the indentation of a left brace on its own line by the value of the variable `c-brace-offset`. For example, with a tab size of eight and default settings for other variables, a `c-brace-offset` of 2 produces:

```
if (a)
{
    b();
}
```

The variable `c-top-braces` controls how much Epsilon indents the braces of the top-level block of a function. By default, Epsilon puts these braces at the left margin. Epsilon indents pre-ANSI K&R-style parameter declarations according to the variable `c-param-decl`. Epsilon indents parts of a top-level structure or union according to `c-top-struct`, and indents continuation lines outside of any function body according to `c-top-contin`. Continuation lines for classes and functions that use C++ inheritance syntax may be indented according to `c-align-inherit`.

Additional C mode indentation variables that may be customized include `c-indent-after-extern-c`, `c-align-break-with-case`, `c-indent-after-namespace`, and `reindent-c-preprocessor-lines`.

By default, the C indenter tries to align continuation lines under parentheses and other syntactic items on prior lines. If Epsilon can't find anything on prior lines to align under, it indents continuation lines two levels more than the original line. (With default settings, Epsilon indents unalignable continuation lines 8 positions to the right of the original line.) Epsilon adds the value of the variable `c-contin-offset` to this

indentation, though. If you want Epsilon to indent unalignable continuation lines ten columns less, set `c-contin-offset` to `-10` (it's `0` by default).

If aligning the continuation line would make it start in a column greater than the value of the variable `c-align-contin-lines` (default `48`), Epsilon won't align the continuation line. It will indent by two levels plus the value of `c-contin-offset`, as described above. A negative value for `c-align-contin-lines` removes the limit entirely. Also see the `c-align-extra-space` variable for an adjustment Epsilon makes for continuation lines that would be indented exactly one level.

As a special case, setting the `c-align-contin-lines` to zero makes Epsilon never try to align continuation lines under syntactic features on prior lines. Epsilon will then indent all continuation lines by one level more than the original line (one extra tab, normally), plus the value of the variable `c-contin-offset`.

If the continuation line contains only a left parenthesis character (ignoring comments), Epsilon can align it with the start of the current statement if you set `c-align-open-paren` nonzero. If the variable is zero, it's aligned like other continuation lines.

You can also have Epsilon use less indentation when a line is very wide. The variable `c-align-contin-max-width` sets a maximum line width for continuation lines, when nonzero. Set it to `-1` to use the current window's width.

When a continuation line is wider than that many columns, the `c-align-contin-max-offset` variable says what to do about it. If greater than zero, Epsilon indents by that amount past the base line (similar to how `c-contin-offset` works). If zero, Epsilon right-aligns the wide line to `c-align-contin-max-width`. If negative, it right-aligns but with that amount of extra space.

These "max" variables, unlike `c-align-contin-lines`, look at the total width of the line, not just the width of its indentation.

C mode also provides special indenting logic for various macros used in Microsoft development environments that function syntactically like braces, such as `BEGIN_ADO_BINDING()`. See the `use-c-macro-rules` variable.

In Objective-C code, Epsilon right-aligns the selectors (argument labels) of multi-line messages, according to the `c-align-selectors` variable.

In C mode, you can use the `find-linked-file` command on `Ctrl-x Ctrl-l` to read the header file included with a `#include` or `#import` statement on the current line, or use the `copy-include-file-name` on `Ctrl-c Alt-i` in a header file to create a suitable `#include` statement. See the `include-directories` variable, and the `mac-framework-dirs` variable for includes that depend on Macintosh framework search paths.

DISABLING C MODE INDENTING

If you prefer manual indenting, various aspects of C mode's automatic indentation can be disabled. If you don't want keys like `#` or `:` or curly braces to reindent the current line, just bind those keys in C mode to `normal-character`. Set `reindent-after-c-yank` and `c-reindent-previous-line` to zero to disable reindenting when yanking, and keep indenting commands from fixing up earlier lines. If you want the `(Enter)` key to go to the next line without indenting, while `Ctrl-j` still does both, you can define a keyboard macro for the former key. Similarly, if you want smart indenting from the `(Tab)` key but a plainer indent from `Ctrl-i`, you can define that by binding `do-c-indent` to the former and one of `indent-previous`, `indent-under`, `indent-like-tab`, or `normal-character` to `Ctrl-i`.

(In a Unix terminal environment, Epsilon can't distinguish keys like `(Enter)` and `(Tab)` from `Ctrl-m` and `Ctrl-i`, respectively, so you'd need to pick different keys.)

Here is an example of the changes to accomplish this.

```
~c-tab "#": normal-character
~c-tab "): normal-character
```

```

~c-tab " ": normal-character
~c-tab "]" : normal-character
~c-tab "{ " : normal-character
~c-tab "} " : normal-character
(set-variable "reindent-after-c-yank" 0)
(set-variable "c-reindent-previous-line" 0)
(define-macro "plain-enter" "C-QC-J")
~c-tab "<EnterKey>": plain-enter
~c-tab "<TabKey>": do-c-indent
~c-tab "C-I": indent-previous

```

Pick the customizations you want, modify them as appropriate, and copy them to your `einit.ecm` customization file (see page 164). Epsilon will begin using the changes the next time it starts up (or use `load-buffer` to load them immediately).

A useful technique when customizing language mode bindings like the above is to run the `list-all` command, then copy the particular lines you want to change into your `einit.ecm` file and modify them. See page 158.

Summary:

	c-mode
C Mode only: <code><Tab></code>	<code>do-c-indent</code>
C Mode only: <code>{</code>	<code>c-open</code>
C Mode only: <code>}</code>	<code>c-close</code>
C Mode only: <code>:</code>	<code>c-colon</code>
C Mode only: <code>#</code>	<code>c-hash-mark</code>
C Mode only: <code>),]</code>	<code>show-matching-delimiter</code>

Other C mode Features

In C mode, the `Alt-⟨Down⟩` and `Alt-⟨Up⟩` keys move to the next or previous `#if/#else/#endif` preprocessor line. When starting from such a line, Epsilon finds the next/previous matching one, skipping over inner nested preprocessor lines. `Alt-]` and `Alt-[` do the same. Press `Alt-i` to display a list of the preprocessor conditionals that are in effect for the current line.

When the cursor is on a brace, bracket, or parenthesis, Epsilon will try to locate its matching brace, bracket, or parenthesis, and highlight them both. If the current character has no match, Epsilon will not highlight it. Similarly, when the cursor is at a preprocessor conditional like `#if`, `#else`, or `#endif`, Epsilon will highlight it and its matching lines. Bits in the variable `auto-show-c-delimiters` enable or modify the rules for both types of highlighting.

C mode colors `#if 0` blocks as comments. The `c-color-preproc-if-pattern` variable determines what counts as an `#if 0` block, so you can make other `#if`-like preprocessor lines also begin comment coloring by customizing this variable.

Press `Alt-'` to display a list of all functions and global variables defined in the current file. You can move to a definition in the list and press `<Enter>` and Epsilon will go to that definition, or press `Ctrl-g` to remain at the starting point. By default, this command skips over external declarations. With a prefix numeric argument, it includes those too. Also see the `list-which-definitions` variable.

Epsilon normally auto-fills text in block comments as you type, breaking overly long lines. See the `c-auto-fill-mode` variable. As with normal auto-fill mode (see page 77), use `Ctrl-x F` to set the right margin for filling. Set the `c-fill-column` variable to change the default right margin in C mode buffers.

Set `fill-c-comment-plain` nonzero if you want block comments to use only spaces instead of a `*` on successive lines.

You can manually refill the current paragraph in a block comment (or in a comment that follows a line of code) by pressing `Alt-q`. If you provide a numeric prefix argument to `Alt-q`, say by typing `Alt-2 Alt-q`, it will fill using the current column as the right margin.

Epsilon's tagging facility isn't specific to C mode, so it's described elsewhere (see page 50). But it's one of Epsilon's most useful software development features, so we mention it here too.

Epsilon displays the name of the current function on the mode line as you move through the buffer. See the `display-definition` variable to disable or customize this. This feature assumes that a right brace character `}` found at the left margin marks the end of a top-level block (such as a function, class, or structure definition). If you must edit code formatted with the `}` of inner blocks also on the left margin, set the `c-quick-scan` variable to zero. (On very large files, this can make updating the function name shown on the mode line much slower.)

Whenever you use the `find-file` command to read in a file with one of the extensions `.c`, `.h`, `.e`, `.y`, `.cpp`, `.cxx`, `.java`, `.inl`, `.hpp`, `.idl`, `.acf`, `.cs`, `.i`, `.ii`, `.m`, `.mi`, `.mm.`, `.mmi`, or `.hxx`, Epsilon automatically enters C mode. See page 84 for information on adding new extensions to this list, or preventing Epsilon from automatically entering C mode. For file names without a suffix, Epsilon examines their contents and guesses whether the file is C++, Perl, some other known type, or unrecognizable.

Summary:	C Mode only: <code>Alt-]</code> , <code>Alt-⟨Down⟩</code>	<code>forward-ifdef</code>
	C Mode only: <code>Alt-[</code> , <code>Alt-⟨Up⟩</code>	<code>backward-ifdef</code>
	C Mode only: <code>Alt-q</code>	<code>fill-comment</code>
	<code>Alt-'</code>	<code>list-definitions</code>
	<code>Alt-i</code>	<code>list-preprocessor-conditionals</code>

4.4.4 Configuration File Mode

Epsilon automatically enters Conf mode when you read a file with an extension of `.conf`, or (under Unix only) when you read a non-binary file in the `/etc` directory. In Conf mode, Epsilon does some generic syntax highlighting, recognizing `#` and `;` as commenting characters, and highlighting `name=value` assignments. It also breaks and fills comments, and provides delimiter highlighting using the `auto-show-conf-delimiters` and `conf-auto-show-delim-chars` variables.

Summary:	<code>conf-mode</code>
----------	------------------------

4.4.5 GAMS Mode

Epsilon automatically enters GAMS mode when you read a file with an extension of `.gms` or `.set`. In addition, if you set the `gams-files` variable nonzero, it recognizes `.inc`, `.map`, and `.dat` extensions. Epsilon also uses GAMS mode for files with an unrecognized extension that start with a `GAMS $title` directive. The GAMS language is used for mathematical programming.

In GAMS mode, Epsilon does syntax highlighting, recognizing GAMS strings and comments. The GAMS language permits a file to define its own additional comment character sequences, besides the standard `*` and `$ontext` and `$offtext`, and Epsilon recognizes most common settings for these.

When the cursor is on a bracket or parenthesis, Epsilon will try to locate its matching bracket or parenthesis, and highlight them both. If the current character has no match, Epsilon will not highlight it. Set the variable `auto-show-gams-delimiters` to zero to disable this feature.

If you use the `compile-buffer` command to compile a GAMS file, Epsilon will automatically look for the `.lst` file produced by GAMS software and translate its format so that the commands `next-error` and `previous-error` work.

Summary:

`gams-mode`

4.4.6 HTML, XML, and CSS Modes

Epsilon automatically enters HTML mode when you read a file with an extension of `.htm`, `.html`, `.shtml`, `.cfml`, `.cfm`, `.htm`, `.asp`, `.asa`, `.htt`, `.jsp`, `.prx`, `.cfc`, `.asx`, `.ascx`, or `.aspx`. It uses XML mode when you read a file with an extension of `.xml`, `.cdf`, `.osd`, `.wml`, `.xsl`, `.xst`, `.xsd`, `.xmp`, `.rdf`, `.svg`, `.rss`, `.xsdconfig`, `.ui`, `.xaml`, `.sgml`, or `.sgm`, or when a file identified as HTML has contents that suggest XML. Everything about HTML mode below also applies to XML mode.

In HTML mode, Epsilon does appropriate syntax highlighting, smart indenting, and brace-matching. The commenting commands and `find-linked-file` work too. If a file contains embedded blocks of CSS or scripting using JavaScript, VBScript, Python, and PHP, Epsilon colors them appropriately, and features like indenting switch to the rules for those languages within such blocks.

The `html-auto-fill-mode` variable controls whether Epsilon automatically breaks long lines as you type. One bit disables this entirely, while others let you customize which lines Epsilon can split (along with the `html-auto-fill-combine` variable). The `html-auto-indent` variable controls when Epsilon indents these new lines (as well as lines you create by pressing `(Enter)`).

For XML mode, Epsilon uses the `xml-auto-fill-mode`, `xml-auto-fill-combine`, and `xml-auto-indent` variables instead.

The `html-indenting-rules` variable controls whether and how Epsilon does smart indenting. The `html-indent` variable sets the width of each level of indentation in HTML text; `xml-indent` is used for XML.

The `html-no-indent-elements` variable lists HTML elements whose contents shouldn't receive additional indentation, and the `html-paragraph-is-container` buffer-specific variable helps Epsilon indent `<p>` tags correctly. Epsilon uses the `html-empty-elements` and `coldfusion-empty-elements` variables to decide which elements never use end tags. The `html-paragraph-elements` variable (`xml-paragraph-elements` in XML) sets which tag names serve as paragraph delimiters for fill-paragraph. The `html-style-rules` variable can be set to require that all empty elements use self-terminating tags.

The `html-reindent-previous-line` variable, or for XML mode, the `xml-reindent-previous-line` variable, controls whether those modes reindent the current line when you press `(Enter)`.

When an HTML or XML file contains embedded scripting, Epsilon must determine its language. If the file itself doesn't specify a language (using a syntax like `<script language=javascript>` or `<?php ?>`, for instance), then Epsilon consults one of several variables to choose a default script language. Each of the variables in figure 4.5 may be set to 1 for Javascript-style coloring, 2 for VBScript-style coloring, 3 for PHP-style coloring, 4 for Python, 5 for CSS, 10 to ignore those block start delimiters, or 0 for plain coloring of the block.

You can disable coloring of embedded scripting in cases where the script language is explicitly specified by setting the `html-prevent-coloring` variable.

Variable	Default	Embedding Syntax
html-asp-coloring	JavaScript	<% %> (in HTML mode)
html-php-coloring	PHP	<? ?> (in HTML mode)
xml-asp-coloring	Ignore	<% %> (in XML mode)
xml-php-coloring	Ignore	<? ?> (in XML mode)
html-other-coloring	JavaScript	<script language= <i>unknown</i> >

Figure 4.5: Variables that control how Epsilon interprets embedded scripting

As you move about in an HTML or XML buffer, on the mode line Epsilon displays the innermost tags in effect for the start of the current line. (The `html-display-nesting-width` variable influences how many tags will be shown.) It shows an open `<` to indicate the line begins inside a tag, or `?` to indicate a syntax problem with the tag such as a missing `>`. Within embedded scripting, it displays the current function's name, or the type of scripting. Set the `html-display-definition` variable to customize this.

When the cursor is on a `<` or `>` character, Epsilon will try to locate its matching `>` or `<` and highlight them both. If the current character has no match, Epsilon will not highlight it. Within a start tag or end tag, Epsilon will use color to highlight it and its matching tag, using a different color to indicate a mismatched tag. Set the variable `auto-show-html-delimiters` to customize this.

Press `Alt-Shift-L` to display a list of all mismatched start or end tags in the current buffer. The list appears in a `grep` buffer. See the `grep-mode` command for navigation details.

When the cursor is at a start tag, you can press `Alt=` to have Epsilon move to its matching end tag, and vice versa.

Press `Alt-Shift-F` to move to the end of the current tag. If the tag is a start tag, Epsilon moves to the end of its matching end tag. Press `Alt-Shift-B` to move to the start of the current tag. On an end tag, it moves to the beginning of its matching start tag. Outside a tag, both commands move by words.

Press `Alt-Shift-D` to delete the current tag, and if it has a matching end or start tag, that tag as well. `Alt-Shift-E` inserts an end tag for the most recent start tag without one.

Mainly useful for XML, the `Alt-Shift-R` key sorts the attributes of the current tag alphabetically by attribute name. With a highlighted region, it sorts the attributes of each tag in the region, then aligns corresponding attributes so they start at the same column in each tag.

Press `Alt-i` to display the element nesting in effect at point. This is similar to the information Epsilon displays in the mode line, but more complete. One difference: while the automatic mode line display shows nesting in effect at the beginning of the current line, and doesn't change as you move around within a line, this command uses the current position within the line.

Epsilon can create an HTML version of syntax-highlighted text that preserves its colors. See the `copy-formatting-as-html` command.

See page 95 for the similar PHP mode. Also see page 130 for information on viewing `http://` URLs with Epsilon.

Epsilon enters CSS mode when you read a file with a `.css` extension. It also uses a flavor of CSS mode when cascading style sheet code is embedded in an HTML file. CSS mode provides syntax highlighting, commenting, smart indenting using the `css-indent` variable, and delimiter highlighting using the `auto-show-css-delimiters` variable.

Summary:

html-mode

	xml-mode
	css-mode
HTML/XML only: Alt-=	html-find-matching-tag
HTML/XML only: Alt-i	html-list-element-nesting
HTML/XML only: Alt-Shift-F	html-forward-tag
HTML/XML only: Alt-Shift-B	html-backward-tag
HTML/XML only: Alt-Shift-D	html-delete-tag
HTML/XML only: Alt-Shift-E	html-close-last-tag
HTML/XML only: Alt-Shift-L	html-list-mismatched-tags
HTML/XML only: Alt-Shift-R	xml-sort-by-attribute-name

4.4.7 Ini File Mode

Epsilon automatically enters Ini mode when you read a file with an extension of .ini or .ecm, and with some files using a .sys or .inf extension. In Ini mode, Epsilon does appropriate syntax highlighting and comment filling (controlled by a bit in the `misc-language-fill-mode` variable).

Summary: ini-mode

4.4.8 Makefile Mode

Epsilon automatically enters Makefile mode when you read a file named makefile (or Makefile, etc.) or with an extension of .mak or .mk. In Makefile mode, Epsilon does appropriate syntax highlighting, and can break and fill comments. The `compile-buffer` command uses the `compile-makefile-cmd` variable in this mode.

Press Alt-i to display a list of the preprocessor conditionals that are in effect for the current line. Epsilon recognizes the preprocessor syntaxes of both Gnu Make and Microsoft Nmake.

Summary: makefile-mode
 Makefile mode only: Alt-i list-make-preprocessor-conditionals

4.4.9 NSIS Mode

Epsilon automatically enters NSIS mode when you read a file with an extension of .nsi, .nsh or .nsdinc. In NSIS mode, Epsilon does appropriate syntax highlighting, recognizing strings, comments, keywords, preprocessor lines, and so forth for files written in the NSIS installer's scripting language. It also breaks and fills comment lines.

When the cursor is on a brace, bracket, or parenthesis, Epsilon will try to locate its matching brace, bracket, or parenthesis, and highlight them both. If the current character has no match, Epsilon will not highlight it. Set the variable `auto-show-nsis-delimiters` to zero to disable this feature.

Summary: nsis-mode

4.4.10 Perl Mode

Epsilon automatically enters Perl mode when you read a file with an extension of `.perl`, `.pm`, `.al`, `.ph`, or `.pl` (or when you read a file with no extension that starts with a `#!` line mentioning Perl). The `compile-buffer` command uses the `compile-perl-cmd` variable in this mode.

Epsilon includes a `perldoc` command that you can use to read Perl documentation. It runs the command of the same name and displays the result in a buffer. You can double-click on a reference to another `perldoc` page, or press `(Enter)` to follow a reference at point, or press `m` to be prompted for another `perldoc` topic name.

Epsilon's syntax highlighting uses the `perl-comment` color for comments and POD documentation, the `perl-function` color for function names, and the `perl-variable` color for variable names.

Epsilon uses the `perl-constant` color for numbers, labels, the simple argument of an angle operator such as `<INPUT>`, names of imported packages, buffer text after `__END__` or `__DATA__`, here documents, format specifications (apart from any variables and comments within), and the operators `my` and `local`.

A here document can indicate that its contents should be syntax highlighted in a different language, by specifying a terminating string with an extension. At the moment the extensions `.tex` and `.html` are recognized. So for example a here document that begins with `<<"end.html"` will be colored as HTML. Epsilon also tries to auto-detect when a here document contains HTML. See the `perl-heredoc-syntax` variable.

Epsilon uses the `perl-string` color for string literals of all types (including regular expression arguments to `s///`, for instance). Interpolated variables and comments are colored appropriately whenever the string's context permits interpolation.

Epsilon uses the `perl-keyword` color for selected Perl operators (mostly those involved in flow control, like `foreach` or `goto`, or with special syntax rules, like `tr` or `format`), and modifiers like `/x` after regular expressions.

Perl mode's automatic indentation features use a modified version of C mode. See page 86 for information on customizing indentation. Perl uses a different set of customization variables whose names all start with `perl-` instead of `c-` but work the same as their C mode cousins. These include `perl-align-contin-lines`, `perl-brace-offset`, `perl-closeback`, `perl-contin-offset`, `perl-label-indent`, `perl-top-braces`, `perl-top-contin`, `perl-top-struct`, and `perl-topindent`. Set `perl-tab-override` if you want Epsilon to assume that tab characters in Perl files aren't always 8 characters wide. Set `perl-indent` if you want to use an indentation in Perl files that's not equal to one tab stop. Set `reindent-perl-comments` to keep indent-region from reindenting comments.

When the cursor is on a brace, bracket, or parenthesis, Epsilon will try to locate its matching brace, bracket, or parenthesis, and highlight them both. If the current character has no match, Epsilon will not highlight it. Set the variable `auto-show-perl-delimiters` to zero to disable this feature.

When you yank blocks of text into a buffer in Perl mode, Epsilon can automatically reindent it. See the variable `reindent-after-perl-yank` to enable this behavior. Some Perl syntax is sensitive to indentation, and Epsilon's indenter may change the indentation, so you should examine yanked text to make sure it hasn't changed.

If you run Perl's debugger inside Epsilon's process buffer under Windows, the following environment variable settings are recommended:

```
set PERLDB_OPTS=ReadLine=0 ornaments=''
set EMACS=yes
```

Summary:

```
perl-mode
perldoc
```

4.4.11 PHP Mode

Epsilon automatically enters PHP mode when you read a file with an extension of .php, .php3, .php4, or .sphp. In PHP mode, Epsilon does appropriate syntax highlighting, tagging, and similar tasks. PHP mode is almost identical to HTML mode. (See page 91.) In both, codes such as `<? ?>` mark PHP scripting, and text outside of these markers is treated as HTML.

When the cursor is on a brace, bracket, or parenthesis in PHP code, Epsilon will try to locate its matching brace, bracket, or parenthesis, and highlight them both. If the current character has no match, Epsilon will not highlight it. Set the variable `auto-show-php-delimiters` to zero to disable this feature.

When you use the `indent-for-comment` command to insert a comment, the `php-comment-style` variable controls which type of comment Epsilon inserts. The value 1 (the default) inserts `#`, the value 2 inserts `//`, and any other value inserts `/*`.

PHP mode's automatic indentation features use a modified version of C mode. See page 86 for information on customizing indentation. PHP uses a different set of customization variables whose names all start with `php-` instead of `c-` but work the same as their C mode cousins. These include `php-align-contin-lines`, `php-brace-offset`, `php-closeback`, `php-contin-offset`, `php-label-indent`, `php-reindent-previous-line`, `php-top-braces`, `php-top-contin`, `php-top-struct`, and `php-topindent`. Set `php-indent` if you want to use an indentation in PHP files that's not equal to one tab stop. The `php-top-level-indent` variable sets the indentation of PHP code outside any function definition.

PHP's syntax highlighting shares its color classes with Perl mode. The color class `perl-comment`, for instance, defines the color of comments in both languages.

A PHP here document can indicate that its contents should be syntax highlighted in a different language, by specifying a terminating string with an extension. At the moment the extensions `.tex` and `.html` are recognized. So for example a here document that begins with `<<"end.html"` will be colored as HTML. Epsilon also tries to auto-detect when a here document contains HTML. See the `php-heredoc-syntax` variable.

Summary:

`php-mode`

4.4.12 PostScript Mode

Epsilon automatically enters PostScript mode when you read a file with an extension of .ps or .eps, or if it contains a PostScript marker on its first line. In PostScript mode, Epsilon does appropriate syntax highlighting, recognizing text strings, comments, and literals like `/Name`. It also breaks and fills comment lines.

When the cursor is on a brace, bracket, or parenthesis, Epsilon will try to locate its matching brace, bracket, or parenthesis, and highlight them both. If the current character has no match, Epsilon will not highlight it. Set the variable `auto-show-postscript-delimiters` to zero to disable this feature.

Summary:

`postscript-mode`

4.4.13 Python Mode

Epsilon automatically enters Python mode when you read a file with an extension of .py or .jy. In Python mode, Epsilon does appropriate syntax highlighting. Tagging, comment filling, and other commenting

commands are also available. Auto-indenting adds an extra level of indentation after a line ending with “:”, a continuation line, or one with an open delimiter, and repeats the previous indentation otherwise.

Set the `python-indent` variable to alter the level of indentation Epsilon uses. Tab widths in Python files are normally set to 8, as required by Python language syntax rules, but you can set the `python-tab-override` variable to change this, or `python-indent-with-tabs` to change whether Python mode uses tabs for indenting, not purely spaces.

Pressing `(Backspace)` to delete a space can delete multiple spaces, as specified by the `python-delete-hacking-tabs` variable.

When the cursor is on a brace, bracket, or parenthesis, Epsilon will try to locate its matching brace, bracket, or parenthesis, and highlight them both. If the current character has no match, Epsilon will not highlight it. Set the variable `auto-show-python-delimiters` to zero to disable this feature.

Set `compile-python-cmd` to modify the command line used by the `compile-buffer` command for Python buffer. Set `python-language-level` to change the list of keywords for syntax highlighting.

Summary: `python-mode`

4.4.14 R Mode

Epsilon automatically enters R mode when you read a file with an extension of `.r`. In R mode, Epsilon does appropriate syntax highlighting, recognizing strings, comments, operators, keywords, functions and other identifiers in files written in the R language. It also breaks and fills comment lines.

When the cursor is on a brace, bracket, or parenthesis, Epsilon will try to locate its matching brace, bracket, or parenthesis, and highlight them both. If the current character has no match, Epsilon will not highlight it. Set the variable `auto-show-rmode-delimiters` to zero to disable this feature.

Summary: `r-mode`

4.4.15 Shell Mode

Epsilon automatically enters shell mode when you read a file with an extension of `.sh`, `.csh`, `.ksh`, `.bash`, `.tcsh`, or `.zsh`, or a name (or extension) of `.bashrc`, `.cshrc`, `.tcshrc`, `.bash_profile`, `.bash_logout`, `.login`, or `.logout`, or when you read a file with no extension that starts with a `#!` line and uses one of these shell names. In Shell mode, Epsilon does appropriate syntax highlighting, recognizing comments, variables and strings.

In Shell mode, Epsilon uses a tab size setting specified by the `shell-tab-override` variable.

When the cursor is on a brace, bracket, or parenthesis, Epsilon will try to locate its matching brace, bracket, or parenthesis, and highlight them both. If the current character has no match, Epsilon will not highlight it. Set the variable `auto-show-shell-delimiters` to zero to disable this feature.

A here document can indicate that its contents should be syntax highlighted in a different language, by specifying a terminating string with an extension. At the moment the extensions `.tex` and `.html` are recognized. So for example a here document that begins with `<<"end.html"` will be colored as HTML. See the `shell-heredoc-syntax` variable.

Summary: `shell-mode`

4.4.16 Tcl Mode

Epsilon automatically enters Tcl mode when you read a file with an extension of .tcl or .ttml. In Tcl mode, Epsilon does appropriate syntax highlighting and smart indenting. Indenting uses the indentation level specified by the `tcl-indent` variable. The mode also breaks and fills comment lines.

When the cursor is on a brace, bracket, or parenthesis, Epsilon will try to locate its matching brace, bracket, or parenthesis, and highlight them both. If the current character has no match, Epsilon will not highlight it. Set the variable `auto-show-tcl-delimiters` to zero to disable this feature.

Summary:

tcl-mode

4.4.17 TeX and LaTeX Modes

Epsilon automatically enters either TeX or LaTeX mode when you read a file with an extension of .tex, .ltx, .sty, or (in most cases) .cls. (By default it enters LaTeX mode, but see the `tex-force-latex` command below.) TeX and LaTeX modes are almost identical, and will be described together.

Keys in TeX/LaTeX mode include Alt-i for italic text, Alt-Shift-I for slanted text, Alt-Shift-T for typewriter, Alt-Shift-B for boldface, Alt-Shift-C for small caps, Alt-Shift-F for a footnote, and Alt-s for a centered line.

Alt-Shift-E prompts for the name of a LaTeX environment, then inserts `\begin{env}` and `\end{env}` lines for the one you select. You can press ? to select an environment from a list. (The list of environments comes from the file `latex.env`, which you can edit by copying it from Epsilon's install directory to your Epsilon customization directory.) Alt-Shift-Z searches backwards for the last `\begin{env}` directive without a matching `\end{env}` directive. Then it inserts the correct `\end{env}` directive at point.

For most of these commands, you can highlight a block of text first and Epsilon will insert formatting commands to make the text italic, slanted, etc. or you can use the command and then type the text to be italic, slanted, etc.

By default, Epsilon inserts the appropriate LaTeX 2_ε/3 command (such as `\textit` for italic text). Set the variable `latex-2e-or-3` to 0 if you want Epsilon to use the LaTeX 2.09 equivalent. (In the case of italic text, this would be `\it`.)

The keys '{' and '\$' insert matched pairs of characters (either {} or \$\$). When you type \ (or \[, TeX/LaTeX mode will insert a matching \) or \], respectively. But if you type '{' just before a non-whitespace character, it inserts only a '{'. This makes it easier to surround existing text with braces.

The keys <Comma> and <Period> remove a preceding italic correction \/, the " key inserts the appropriate kind of doublequote sequence like ‘ ‘ or ' ', and Alt-" inserts an actual " character.

Some TeX mode commands are slightly different in LaTeX than in pure TeX. Set `tex-force-latex` to 1 if all your documents are LaTeX, 0 if all your documents are TeX, or 2 if Epsilon should determine this on a document-by-document basis. In that case, Epsilon will assume a document is LaTeX if it contains a `\begin{document}` statement or if it's in a file with an .ltx, .sty, or .cls extension. By default, Epsilon assumes all documents use LaTeX.

When the cursor is on a delimiter character like (,), {, }, [, or], Epsilon will try to locate its matching character and highlight them both. If the current character has no match, Epsilon will not highlight it. Set the variable `auto-show-tex-delimiters` to zero to disable this feature.

Set the variable `tex-look-back` to a bigger number if you want TeX mode to more accurately syntax highlight very large paragraphs but be slower, or a smaller number if you want recoloring to be faster but

perhaps miscolor large paragraphs. You can customize syntax highlighting using the variables `latex-display-math-env-pat`, `latex-math-env-pat`, and `latex-non-text-argument`.

You can use the `list-definitions` command to see a list of LaTeX labels in the current file and move to one. The tagging commands (see page 50) also work on labels. See the `latex-tag-keywords` variable if you want to make these work on cite tags too, or make other tagging customizations.

In LaTeX mode, the spell checker uses the `latex-spell-options` variable. Also see the `latex-non-text-argument` variable to control how the spell checker treats the parameter of LaTeX commands like `\begin` that can take keywords. In TeX mode, the spell checker uses the `tex-spell-options` variable.

RUNNING TEX

The `compile-buffer` command (normally on Alt-F3) uses the `compile-tex-cmd` variable in TeX mode and the `compile-latex-cmd` variable in LaTeX mode. (We'll just say "TeX" from now on, but we mean either TeX or LaTeX or any comparable program like pdfLaTeX or XeTeX.) You may need to set these if the version of TeX you use takes some different flags. By default Epsilon uses flags that keep TeX from prompting for input on errors, display error messages in a format other programs can read, and enable SyncTeX and source specials for integrating a viewer with Epsilon. You could also use a definition that runs TeX and then displays the resulting output file, but the details depend on which viewer you want to use.

If your TeX system uses a compatible DVI previewer, or you have a PDF viewer that works with the SyncTeX system, then you can configure your viewer and Epsilon to work together, jumping from any line in your `.tex` source file to the resulting TeX output in your viewer (forward sync), or jumping from a spot in your viewer back to the `.tex` code that produced it (reverse sync). This requires some setup. We'll describe reverse sync first.

For SyncTeX, add the flag `-synctex=1` to your TeX command line (by adding it to `compile-latex-cmd`, for example), or put `\synctex=1` in your `.tex` file. Then install a PDF viewer that understands SyncTeX files. On Windows, SumatraPDF can do this. In its configuration file, use these settings:

```
EnableTeXEnhancements = true
InverseSearchCmdLine = "C:\Program Files\Eps14\bin\epsilon.exe" -add +%l "%f"
```

The above setting lets you double-click on text in your PDF and have Epsilon display the corresponding TeX file and line that produced it.

Similarly, to link a DVI previewer to Epsilon, you must first instruct TeX or LaTeX to include "source specials" in your DVI file; these let your DVI viewer know which parts of the DVI file correspond to particular `.tex` file names and line numbers. Most versions of TeX understand flags like `--src-specials` or `-src` to do this. With others, your TeX source file can input `srcex.sty` (or `srcltx.sty`) to include this information.

Once you've set up TeX to use source specials, as above, you can configure your DVI viewer to run Epsilon, showing the source file and line corresponding to a certain spot in your DVI file. The details depend on your DVI viewer, but a command line like `epsilon -add +%l "%f"` is typical.

For yap, use the View/Options/Inverse Search command to set the command line as above. You may need to include the full path to Epsilon. Double-click in yap to have Epsilon display the source file that generated the DVI text you're viewing.

For xdvi, you can run it with its `-editor` flag, or set the `XEDITOR` environment variable, so it contains `epsilon -add +%l:%c %f`. Ctrl-click in xdvi to have Epsilon display the corresponding source file.

In Epsilon, forward sync (jumping from your `.tex` source file to the corresponding TeX output in your DVI or PDF viewer) uses the `jump-to-dvi` command on Alt-Shift-J. Or you can hold down Ctrl and left-click

the mouse on your TeX source. First, enable source specials or SyncTeX support, respectively, as explained above. Then you'll need to configure Epsilon. Set the `jump-to-dvi-extension` variable to either `".dvi"` or `".pdf"`. Then set the `jump-to-dvi-command` variable. It contains the command line to run your DVI or PDF viewer and instruct it to go to a particular TeX source file and line within the output file. For `xdvi`, use `xdvi -sourceposition %1%f %d` or `xdvi -sourceposition %1%b.tex %d`. For `yap`, use `yap -1 -s %1%f %d`. For `SumatraPDF`, use `SumatraPDF -forward-search %f %1 %d`. You might need to add the full path, so a complete setting in your `einitt.ecm` file might look like this:

```
+jump-to-dvi-command char array variable: 4096
  "\"C:\\Program Files\\SumatraPDF\\SumatraPDF.exe\" -forward-search %f %1 %d"
+jump-to-dvi-extension char array variable: 4096 ".pdf"
```

(We've split the first line above because it's so long, but it's really one long line.)

Once both those Epsilon variables have been set appropriately, you can use TeX mode's `jump-to-dvi` command via `Alt-Shift-J` or `Ctrl-left-click` to view the DVI or PDF output resulting from the current line of TeX or LaTeX.

(Under Windows, Epsilon can also communicate with Y&Y's old "DVIWindo" previewer, version 2.1.4 and later, using DDE. For this option, `jump-to-dvi-command` must be empty. When `jump-to-dvi` uses DVIWindo, it doesn't prompt; instead, it assumes that if the current TeX source file has no corresponding DVI file, the DVI file is already loaded in DVIWindo.)

For documents made from multiple TeX files, Epsilon can't always determine the ultimate DVI or PDF file name by examining one of the component TeX files. You can add a line like `"% mainfile: somefilename.tex"` or `"% !TeX root = somefilename.tex"` within 3 lines of a TeX file's start or end, and Epsilon will assume that `somefilename.dvi` (or `somefilename.pdf`, based on the `jump-to-dvi-extension` variable) is the file your viewer should open. If your document doesn't include a `mainfile` declaration, Epsilon will look for a corresponding `.dvi` or `.pdf` file with the same base name. Otherwise, it prompts for the `.dvi` (or `.pdf`) file name. It uses that file name from then on, even if some other TeX file has a `mainfile` declaration or corresponding `.dvi` or `.pdf` file.

If this command displays the wrong DVI or PDF file, run `jump-to-dvi` with a numeric prefix argument, which forces it to ask for the correct file name again. An empty response makes Epsilon return to checking separately for each `.tex` file, as above.

Summary:	<code>Alt-i</code>	<code>tex-italic</code>
	<code>Alt-Shift-I</code>	<code>tex-slant</code>
	<code>Alt-Shift-T</code>	<code>tex-typewriter</code>
	<code>Alt-Shift-B</code>	<code>tex-boldface</code>
	<code>Alt-Shift-C</code>	<code>tex-small-caps</code>
	<code>Alt-Shift-F</code>	<code>tex-footnote</code>
	<code>Alt-s</code>	<code>tex-center-line</code>
	<code>Alt-Shift-E</code>	<code>tex-environment</code>
	<code>Alt-Shift-Z</code>	<code>tex-close-environment</code>
	<code>{</code>	<code>tex-left-brace</code>
	<code>\$</code>	<code>tex-math-escape</code>
	<code><Comma></code> , <code><Period></code>	<code>tex-rm-correction</code>
	<code>"</code>	<code>tex-quote</code>
	<code>Alt-"</code>	<code>tex-force-quote</code>
	<code>\(</code>	<code>tex-inline-math</code>
	<code>\[</code>	<code>tex-display-math</code>

	<code>tex-mode</code>
	<code>latex-mode</code>
<code>Alt-Shift-J, Ctrl-M-(Left)</code>	<code>jump-to-dvi</code>

4.4.18 VHDL Mode

Epsilon automatically enters VHDL mode when you read a file with an extension of `.vhd` or `.vhdl`. In VHDL mode, Epsilon does appropriate syntax highlighting. It also breaks and fills comment lines.

When the cursor is on a parenthesis, Epsilon will try to locate its matching parenthesis, and highlight them both. If the current character has no match, Epsilon will not highlight it. Set the variable `auto-show-vhdl-delimiters` to zero to disable this feature.

Summary:	<code>vhdl-mode</code>
----------	------------------------

4.4.19 Visual Basic Mode

Epsilon automatically enters Visual Basic mode when you read a file with an extension of `.vb`, `.bas`, `.frm`, `.vbs`, `.ctl`, or `.dsr` (plus certain `.cls` files as well). In Visual Basic mode, Epsilon does appropriate syntax highlighting, smart indenting, tagging, and comment filling.

When the cursor is on a brace, bracket, or parenthesis, Epsilon will try to locate its matching brace, bracket, or parenthesis, and highlight them both. If the current character has no match, Epsilon will not highlight it. Set the variable `auto-show-vbasic-delimiters` to zero to disable this feature.

Set the `vbasic-indent` variable to alter the level of indentation Epsilon uses. Set `vbasic-indent-subroutines` to change Epsilon's indenting style. Set the `vbasic-indent-with-tabs` variable nonzero if you want Epsilon to indent using a mix of tab characters and spaces, instead of just spaces.

When you yank blocks of text into a buffer in Visual Basic mode, Epsilon can automatically reindent it. See the variable `reindent-after-vbasic-yank` to enable this behavior. The `vbasic-reindent-previous-line` variable controls whether pressing `(Enter)` to insert and indent a new line also reindents the existing one.

The `vbasic-language-level` variable lets you customize which Visual Basic keywords Epsilon colors.

Summary:	<code>vbasic-mode</code>
----------	--------------------------

4.5 More Programming Features

Epsilon has a number of features that are useful when programming, but work similarly regardless of the programming language. These are described in the following sections. Also see the language-specific commands described in previous sections, the tagging commands on page 50, and the `align-region` command described on page 80.

4.5.1 Navigating in Source Code

In most language modes, you can press `Alt-'` to display a list of all functions and global variables defined in the current file. You can move to a definition in the list and press `<Enter>` and Epsilon will go to that definition, or press `Ctrl-g` to remain at the starting point.

By default, this command skips over external declarations. With a prefix numeric argument, it includes those too (if the current language has such a notion and the mode supports this). The `list-which-definitions` variable lets you customize which types of definitions are shown (in those modes that can show more than one type). The `list-definitions-live-update` variable lets you keep Epsilon from repositioning in the source file window as you navigate in the definitions window or include source file line numbers in the window title.

The `goto-next-declaration` command on `Ctrl-c Ctrl-n` moves to the next global declaration or definition of a function or variable in the current file, and the `goto-previous-declaration` command on `Ctrl-c Ctrl-p` move to the previous one. The `goto-next-definition` and `goto-previous-definition` commands are similar but move by definitions, ignoring declarations.

Also see Epsilon's tagging features on page 50, and its interface to Microsoft's source code browsing database on page 52.

Summary:	<code>Alt-'</code>	<code>list-definitions</code>
	<code>Ctrl-c Ctrl-n</code>	<code>goto-next-declaration</code>
	<code>Ctrl-c Ctrl-p</code>	<code>goto-previous-declaration</code>
		<code>goto-next-definition</code>
		<code>goto-previous-definition</code>

4.5.2 Pulling Words

The `pull-word` command bound to the `Ctrl-⟨Up⟩` key (as well as the `F3` key) scans the buffer before point, and copies the previous word to the location at point. If you type the key again, it pulls in the word before that, and so forth. Whenever Epsilon pulls in a word, it replaces any previously pulled-in word. If you like the word that has been pulled in, you do not need to do anything special to accept it—Epsilon resumes normal editing when you type any key except for the few special keys reserved by this command. You can type `Ctrl-⟨Down⟩` (the `pull-word-fwd` command) to go in the other direction. Type `Ctrl-g` to erase the pulled-in word and abort this command.

If a portion of a word immediately precedes point, that subword becomes a filter for pulled-in words. For example, suppose you start to type a word that begins `WM`, then you notice that the word `WM_QUERYENDSESSION` appears a few lines above. Just type `Ctrl-⟨Up⟩` and Epsilon fills in the rest of this word.

The command provides various visual clues that tell you exactly from which point in the buffer Epsilon is pulling in the word. If the source is close enough to be visible in the window, it is simply highlighted. If the pulled-in word comes from farther away, Epsilon shows the context in the echo area, or in a context window that it pops up (out of the way of your typing).

When there are no more matches before point in the current buffer, Epsilon loads a tag file (see page 50) and looks for matches there. The `pull-word-from-tags` variable controls this behavior. In this way, you can complete on any tagged identifier by typing part of it and pressing `F3` or `Ctrl-⟨Up⟩` until the identifier you want appears.

You can also pull words from the buffer at most prompts. For instance, you can retrieve a long file name that appears in the buffer into a `find-file` prompt, or look for other instances of an identifier in a program

without typing the whole identifier. Type the first few characters at a search or `grep` prompt and press `F3` or `Ctrl-⟨Up⟩` to pull the rest.

Summary:	<code>Ctrl-⟨Up⟩, F3</code>	<code>pull-word</code>
	<code>Ctrl-⟨Down⟩</code>	<code>pull-word-fwd</code>

4.5.3 Accessing Help

This section describes how Epsilon can help you access programming language and library documentation. See page 37 for directions on obtaining help on Epsilon itself.

To get help on the word at point, press `Shift-F1` to run the `context-help` command. It provides help on the identifier at point, selecting the appropriate type of help based on the current mode. See page 102 for details. In some modes, it uses commands explained in this section.

Epsilon for Unix provides a `man` command for reading man pages. At its prompt, type anything you would normally type to the `man` command, such as `-k open` to get a list of man pages related to the keyword “open”. If you don’t use any flags or section names, Epsilon will provide completion on available topics. For example, type “?” to see all man page topics available. Within man page output, you can double-click on a reference to another man page, such as `echo(1)`, or press `⟨Enter⟩` to follow it, or press `m` to be prompted for another man page topic. The `man` command also works under Windows using the Cygwin environment, if its `man` program is installed.

The `search-man-pages` command generates a list of man pages that contain some specified text, putting its results in the `grep` buffer. (See page 48.) It first prompts for the search string. You can use `Ctrl-t`, `Ctrl-c`, or `Ctrl-w` to toggle regular expression, case folding, or word searching behavior, as with `grep` and other searching commands.

Then it asks if you want to restrict searching to particular man page sections, such as `1` for commands or `3` for subroutines. Use `*` to search all sections. Finally, it asks if you want to restrict the search to man page entries matching a certain file pattern, such as `*file*` to search only pages whose names contain “file”.

For speed reasons, it searches each man page without processing it through the `man` command, searching the man page in source format. By default, it shows only the first match in each page; set the `search-man-pages-shows-all` variable to see all matches. The result appears in the `grep` buffer; when you view a match from there, Epsilon will then use the `man` command to display its processed form.

Epsilon also includes a `perldoc` command for reading Perl documentation. Just like `man`, it works by running an external program, in this case the `perldoc` program that comes with Perl. Run `perldoc` on the topic `perldoc` to see the flags you can use with it, such as `-f` to locate the documentation for a specific Perl function.

Summary:	<code>man</code>
	<code>perldoc</code>
	<code>search-man-pages</code>

4.5.4 Context-Sensitive Help

Pressing `Shift-F1` provides help on the keyword at point, using a help rule specific to the current mode. Modes can have rules that consult compiler documentation, man or Info pages, or search on the web at an API reference site. You can change the rules by setting variables.

Each mode has an associated variable whose name is formed by adding the mode name to the end of `context-help-rule-`. (Characters in the mode name that aren't valid in EEL identifier names are removed first; the rest are changed to lowercase.) So HTML mode, for instance, uses a rule variable named `context-help-rule-html`.

A mode can define separate rules for Windows and Unix. If a mode named `m` defines a variable `context-help-rule-m-windows` or `context-help-rule-m-unix`, Epsilon will use it instead of the usual variable on that platform.

If there's no variable for the current mode, Epsilon uses the rule in the `context-help-default-rule` variable.

Instead of, or in addition to, providing a rule variable, a mode can define an EEL function to force a different rule variable. Epsilon calls a function named `context_help_override_modename` for a mode named *modename*, if one is defined. That function must fill in its single `char *` parameter with a replacement mode name; then Epsilon will look up the corresponding variable, as above. C mode uses this to provide a different help rule for EEL buffers than for Java or C++ buffers, even though they all use C mode. HTML mode uses this to provide a different help rule within embedded JavaScript (or other scripting languages) than it does in plain HTML text.

Each rule variable consists of a rule type, which is a single punctuation character, followed by a parameter such as a file name or URL. Here are the rule types:

>infofiles looks up the keyword at point in the index of the Info file specified by its parameter. You can list multiple comma-separated Info file names, and Epsilon will use the first one that contains the keyword in its index.

+func runs the specified EEL function *func*. See below for useful rule functions.

!cmdline runs an external program, using the specified command line. It appends the keyword at point to the end of the command line. But if the command line contains a `*` character, Epsilon instead substitutes the keyword at that position. If the program file name contains spaces, quote it with `"`. A `&` character at the end of the command line makes the command run asynchronously; use it to run a Windows GUI program. Without `&`, Epsilon will run the program synchronously, collect its output in a buffer, and display any output it produces.

=url runs a web browser on the specified URL. A `*` character in the URL substitutes the keyword at point at that position; otherwise it's added to the end of the URL.

&site runs a web browser using the URL of a search engine, telling it to perform a site-specific search for the keyword at point. The site name follows the `&` character. The URL is constructed by taking the value of the `context-help-site-search` variable and replacing the `@` in it with the given site name. You can set that variable to change your preferred search provider. The resulting URL is then interpreted as with `=` above.

@site runs a web browser using the URL of a search engine, telling it to perform a site-specific search for the keyword at point, and go to the first match. The site name follows the `@` character. The URL is constructed by taking the value of the `context-help-site-search-lucky` variable and replacing the `@` in it with the given site name. You can set that variable to change your preferred search provider. The resulting URL is then interpreted as with `=` above.

\$helpfile looks up the keyword at point in a `.hlp` file. This rule type is only available under Windows, and only works if you've installed WinHelp (an optional feature through Windows 8.1 only).

Here are some functions designed to be used in a context rule. To use one in a rule, put a `+` before its name in the rule variable.

The `context_help_man()` subroutine runs Epsilon's `man` command on the word at point. (The `_` and `-` characters are interchangeable in these rules, so you can set a rule to `+context-help-man`, for instance, to use this subroutine.)

The `context_help_perldoc()` subroutine runs Epsilon's `perldoc` command on the word at point.

Under Windows, you can set a mode to use help provided by Microsoft's older development environments in MS-Help2 format, by setting its rule to call the `context_help_windows_compilers()` subroutine. (Microsoft no longer uses this help format in their development environments.) The subroutine tries to locate the appropriate compiler help system automatically. With certain help systems that have multiple help collections, it uses the `mshelp2-collection` variable.

Summary: Shift-F1, F1 t context-help

4.5.5 Commenting Commands

The `Alt-;` command creates a comment on the current line, using the commenting style of the current language mode. The comment begins at the column specified by the `comment-column` variable (by default 40). (However, if the comment is the first thing on the line and `indent-comment-as-code` is nonzero, it indents to the column specified by the buffer's language-specific indentation function.) If the line already has a comment, this command moves the comment to the comment column. (Also see the `align-region` command described on page 80 to align all comments in a region to the same column.)

With a numeric argument, `Alt-;` searches for the next comment in the buffer and goes to its start. With a negative argument, `Alt-;` searches backwards for a comment. Press `Alt-;` again to reindent the comment.

By default (and in modes that don't specify a commenting style), comments begin with the `;` character and continue to the end of the line. C mode recognizes both old-style `/* */` comments, and the newer C++-style comments `//`, and by default creates the latter. Set the variable `new-c-comments` to 0 if you want `Alt-;` to create old-style comments.

The `Ctrl-x ;` command sets future comments to begin at the current column. With a positive argument, it sets the comment column based on the indentation of the previous comment in the buffer. If the current line has a comment, this command reindents it.

With a negative argument (as in `Alt-(Minus) Ctrl-x ;`), the `Ctrl-x ;` command doesn't change the comment column at all. Instead, it kills any comment on the current line. The command saves the comment in a kill buffer.

You can comment out a region of text by pressing `Ctrl-c Ctrl-r`. Epsilon adds a comment delimiter to the start of each line in the highlighted region between point and mark (or only the current line, if no region has been highlighted). With a numeric argument, as in `Ctrl-u Ctrl-c Ctrl-r`, Epsilon removes such a comment delimiter from each line. Also see the `comment-region-options` variable.

The comment commands look for comments using regular expression patterns (see page 66) contained in the buffer-specific variables `comment-pattern` (which should match the whole comment) and `comment-start` (which should match the sequence that begins a comment, like `/*`). When creating a comment, it inserts the contents of the buffer-specific variables `comment-begin` and `comment-end` around the new comment. When Epsilon puts a buffer in C mode, it decides how to set these variables based on the `new-c-comments` variable.

In certain modes, including C and Perl modes, Epsilon normally auto-fills text in block comments as you type, breaking overly long lines. See the `c-auto-fill-mode` variable for C and Perl modes, `tex-auto-fill-mode` for TeX, `html-auto-fill-mode` for HTML, `xml-auto-fill-mode` for XML, and `misc-language-fill-mode` for Makefile, VHDL, Visual Basic, Python, PostScript, Conf, Ini, and

many other modes. As with normal auto-fill mode (see page 77), use Ctrl-x f to set the right margin for filling. Set the `c-fill-column` variable to change the default right margin in C and Perl mode buffers; `margin-right` in other modes.

You can manually refill the current paragraph in a block comment by pressing Alt-q. If you provide a numeric prefix argument to Alt-q, say by typing Alt-2 Alt-q, it will fill using the current column as the right margin. By default, Epsilon doesn't apply auto-filling to a comment line that also contains non-comment text (such as a C statement with a comment after it on the same line). Use Alt-q to break such lines.

The `auto-fill-comment-rules` variable lets you customize certain aspects of Epsilon's behavior when breaking and filling comment lines.

Summary:	Alt-;	indent-for-comment
	Ctrl-x ;	set-comment-column
	Alt-(Minus) Ctrl-x ;	kill-comment
	Ctrl-c Ctrl-r	comment-region

4.6 Fixing Mistakes

4.6.1 Undoing

The undo command on F9 undoes the last command, restoring the previous contents of the buffer, or moving point to its position, as if you hadn't done the last command. If you press F9 again, Epsilon will undo the command before that, and so forth.

For convenience, when typing text Epsilon treats each word you type as a single command, rather than treating each character as its own command. For example, if you typed the previous paragraph and pressed undo, Epsilon would remove the text “forth.”. If you pressed undo again, Epsilon would remove “so ”.

Epsilon's undo mechanism considers each subcommand of a complicated command such as query-replace a separate command. For example, suppose you do a query-replace, and one-by-one replace ten occurrences of a string. The undo command would then reverse the replacements one at a time.

Epsilon remembers changes to each buffer separately. Say you changed buffer 1, then changed buffer 2, then returned to buffer 1. Undoing now would undo the last change you made to buffer 1, leaving buffer 2 alone. If you switched to buffer 2 and invoked undo, Epsilon would then undo changes to that buffer.

The redo command on F10 puts your changes back in (it undoes the last undo). If you press undo five times, then press redo four times, the buffer would appear the same as if you pressed undo only once.

You can move back and forth undoing and redoing in this way. However, if you invoke a command (other than undo or redo) that either changes the buffer or moves point, you can not redo any commands undone immediately before that command. For example, if you type “one two three”, undo the “three”, and type “four” instead, Epsilon will behave as if you had typed “one two four” all along, and will let you undo only that.

The commands `undo-changes` and `redo-changes` work like undo and redo, except they will automatically undo or redo all changes to the buffer that involve only movements of point, and stop just before a change of actual buffer contents.

For example, when you invoke `undo-changes`, it performs an undo, then continues to undo changes that involve only movements of point. The `undo-changes` command will either undo a single buffer modification (as opposed to movement of point), as a plain undo command would, or a whole series of movement commands at once. It doesn't undo any movement commands after undoing a buffer modification, only after undoing other movement commands. The `redo-changes` command works similarly.

The Ctrl-F9 key runs undo-changes, and the Ctrl-F10 key runs redo-changes.

The commands undo-by-commands and redo-by-commands are another alternative; they try to group undo operations on a command-by-command basis.

Use the undo-movements command on Ctrl-F11 to move to the location of previous editing operations. It uses the same undo information, but without making any changes to the buffer. The redo-movements command on Ctrl-F12 goes in the opposite direction, toward more recent editing locations.

The buffer-specific variable undo-size determines, in part, how many commands Epsilon can remember. For example, if undo-size has the value 500,000 (the default), Epsilon will save at most 500,000 characters of deleted or changed text for each buffer. Each buffer may have its own value for this variable. Epsilon also places an internal limit on the number of commands, related to command complexity. Epsilon can typically remember about 10,000 simple commands (ignoring any limit imposed by undo-size) but more complicated commands make the number smaller.

Summary:	F9, Ctrl-x u	undo
	F10, Ctrl-x r	redo
	Ctrl-F9, Ctrl-x Ctrl-u	undo-changes
	Ctrl-F10, Ctrl-x Ctrl-r	redo-changes
	Ctrl-F11	undo-movements
	Ctrl-F12	redo-movements
		undo-by-commands
		redo-by-commands

4.6.2 Interrupting a Command

You can interrupt a command by pressing Ctrl-g, the default *abort key*. For example, you can use Ctrl-g to stop an incremental search on a very long file if you don't feel like waiting. You can set the abort key with the set-abort-key command. If you interrupt Epsilon while reading a file from disk or writing a file to disk, it will ask you whether you want to abort or continue. Typing the abort key also cancels any currently executing keyboard macros. Aborting normally only works when a command checks for it.

Summary:	Ctrl-g	abort
		set-abort-key

4.7 The Screen

4.7.1 Display Commands

The Ctrl-l (letter L) command causes Epsilon to center point in the window. If you give a numeric argument to Ctrl-l, Epsilon makes the current line appear on that line of the window. For instance, give a numeric argument of zero to make the current line appear on the topmost line of the window. (The line-to-top command is another way to do this.) If you give a numeric argument greater than the number of lines the window occupies, Epsilon will position the current line at the bottom of the window. (The line-to-bottom command is another way to do this.) When repeated, the Ctrl-l command also completely refreshes the screen. If some other program has written text on the screen, or something has happened to garble the screen, use this command to refresh it.

The Alt-⟨Comma⟩ and Alt-⟨Period⟩ commands move point to the first and last positions displayed on the window, respectively.

The Ctrl-z and Alt-z commands scroll the text in the window up or down, respectively, by one line. These scrolling commands will move point as necessary so that point remains visible in the window.

The Ctrl-v and Alt-v commands scroll the text of the window up or down, respectively, by several lines fewer than the size of the window. These commands move point to the center line of the window.

You can control the exact amount of overlap between the original window of text and the new window with the `window-overlap` variable. A positive value for this variable means to use that number of screen lines of overlap between one window of text and the next (or previous). A negative value for `window-overlap` represents a percentage of overlap, instead of the number of screen lines. For example, the default value for `window-overlap` of 2 means to use 2 lines of overlap. A value of -25 for `window-overlap` means to overlap by 25%.

You can change how Epsilon pages through a file by setting the variable `paging-centers-window`. Epsilon normally positions the cursor on the center line of the window as you move from page to page. Set this variable to zero if you want Epsilon to try to keep the cursor on the same screen line as it pages.

The goto-line command on Ctrl-x g prompts for a line number and then goes to the beginning of that line in the current buffer. If you prefix a numeric argument, Epsilon will use that as the line number. Use the format 10:20 to include a column specification; that one goes to line 10, column number 20. Or use a percent character to indicate a buffer percentage: 25% goes to a line 25% of the way through the buffer. Or use the format p123 to go to a particular buffer offset, counting by characters.

The Ctrl-x L command shows the number of lines in the buffer and the number of the line containing point. It also shows the number of bytes the file would occupy if written to disk. This can differ from the size of the buffer, because the latter counts each line separator as a single character. Such characters require two bytes when written to disk in the format used in Windows and DOS, however. See page 122 for information on how Epsilon translates line separator characters. (The character count shown does not take into account any Unicode encoding that may happen when you save the file, so for such files it represents the number of characters that will be saved, which will vary from the number of bytes on disk.)

See the `mode-format` variable if you want to change how Epsilon displays the current line or column number in the mode line at all times, or `draw-line-numbers` if you want Epsilon to display each line's number in a column to its left.

The Ctrl-x = command displays in the echo area information pertaining to point. It shows the size of the buffer, the character position in the buffer corresponding to point, that character's column, and the value of that character in decimal, hex, and "normal" character representation, as well as the character's name for non-ASCII characters.

The count-words command displays the number of words in the current buffer. Highlight a region first and it will count only words in the region. With a numeric argument, it prompts for a search pattern and then counts the number of instances of that pattern.

Summary:	Ctrl-l	center-window
	Ctrl-v, ⟨PgDn⟩	next-page
	Alt-v, ⟨PgUp⟩	previous-page
	Ctrl-z	scroll-up
	Alt-z	scroll-down
	⟨Home⟩, Alt-⟨Comma⟩	beginning-of-window
	⟨End⟩, Alt-⟨Period⟩	end-of-window
		line-to-top
		line-to-bottom

Ctrl-x =	show-point
Ctrl-x L	count-lines
Ctrl-x g	goto-line
Ctrl-c Ctrl-w	count-words

4.7.2 Horizontal Scrolling

The Alt-{ and Alt-} commands scroll the text in the window to the left or right, respectively, by one column.

The Alt-{ and Alt-} commands also control how Epsilon displays long lines to you. Epsilon can, for display purposes, wrap long lines to the next line. Epsilon indicates a wrapped line by displaying a special continuation character where it broke the line for display purposes. But by default Epsilon displays long lines by simply scrolling them off the display. To switch from scrolling long lines to wrapping long lines, use the Alt-} command to scroll to the right, past the end. Epsilon will then wrap long lines.

Similarly, to switch from wrapping long lines to scrolling long lines, press the Alt-{ key. Subsequent use of the Alt-{ command will then scroll the text in the window to the left, as explained above. Whenever Epsilon changes from one display scheme to the other, it indicates the change in the echo area. If, due to scrolling, some of a buffer's contents would appear past the left edge of the screen, the mode line displays "<number" to indicate the number of columns hidden to the left.

You can also use the change-line-wrapping command to set whether Epsilon wraps long lines in the current window, or horizontally scrolls across them.

If you want Epsilon to always wrap long lines, set the default value of the window-specific variable `display-column` to -1 using the `set-variable` command on F8. See page 162 for info on making such customizations permanent.

In a dialog, another way to handle lines that are too long to fit in a window is to resize the dialog by moving its borders. Most dialogs in Epsilon for Windows are resizable, and Epsilon will remember the new size from session to session.

The Alt-PageUp and Alt-PageDown keys scroll horizontally, like Ctrl-v and Alt-v. More precisely, they move the point left or right on the current line by about half the width of the current window, then reposition the window so the point is visible. The command `jump-to-column` on Alt-g prompts for a column number, then goes to the specified column.

Summary:	Alt-{	scroll-left
	Alt-}	scroll-right
		change-line-wrapping
	Alt-⟨PageUp⟩	page-left
	Alt-⟨PageDown⟩	page-right
	Alt-g	jump-to-column

4.7.3 Windows

Epsilon has quite a few commands to deal with creating, changing, and moving windows. Changing the size or number of the windows never affects the buffers they display.

Normally, each buffer has a single point, but this can prove inconvenient when a buffer appears in more than one window. For this reason, Epsilon associates a point with each window in that case. Consequently, you can look at different parts of the same buffer by having the same buffer displayed in different windows and moving around independently in each of them.

Creating Windows

The `Ctrl-x 2` command splits the current window into two windows, one on top of the other, each about half as large. Each window displays the same buffer that the original did. This command will only split the window if each new window would occupy at least 1 screen line, not counting the mode line. To edit another file in a new window, first use `Ctrl-x 2`, then use one of the file commands described on page 118.

The `Ctrl-x 5` command works similarly, but splits the current window so that the two child windows appear side by side, instead of stacked. This command will only split the window if each new window would occupy at least 1 column. Since this typically results in narrow windows, the `Ctrl-x 5` command also sets up the windows to scroll long lines, as described on page 108. See the `wrap-split-vertically` variable to control this.

When you display the same buffer in several narrow windows side by side, follow mode can be useful. It operates when the same buffer is displayed in adjacent windows, by linking the windows together so scrolling and other movement in one is immediately reflected in the others. The `follow-mode` command toggles this mode for the current buffer. The `follow-mode-overlap` variable controls how much the window text overlaps.

Summary:	<code>Ctrl-x 2</code>	<code>split-window</code>
	<code>Ctrl-x 5</code>	<code>split-window-vertically</code>
		<code>follow-mode</code>

Removing Windows

To get rid of the current window, use the `Ctrl-x 0` command. If the previous window can move into the deleted window's space, it does. Otherwise, the next window expands into the deleted window's space.

The `Ctrl-x 1` command makes the current window occupy the entire screen, deleting all the other windows. The `Ctrl-x z` command operates like `Ctrl-x 1`, except that it also remembers the current window configuration. Later, if you type `Ctrl-x z` again, the command restores the saved window configuration.

Summary:	<code>Ctrl-x 0</code> , <code>Ctrl-x Ctrl-d</code>	<code>kill-window</code>
	<code>Ctrl-x 1</code>	<code>one-window</code>
	<code>Ctrl-x z</code>	<code>zoom-window</code>

Selecting Windows

The `Ctrl-x n` key moves to the next window, wrapping around to the first window if invoked from the last window. The `Ctrl-x p` key does the reverse: it moves to the previous window, wrapping around to the last window if invoked from the first window.

You can think of the window order as the position of a window in a list of windows. Initially only one window appears in the list. When you split a window, the two child windows replace it in the list. The top or left window comes before the bottom or right window. When you delete a window, that window leaves the list.

You can also change windows with the `move-to-window` command. It takes a cue from the last key in the sequence used to invoke it, and moves to a window in the direction indicated by the key. If you invoke the command with `Ctrl-x <Right>`, for example, the window to the right of the cursor becomes the new current window. The `Ctrl-x <Left>` key moves left, `Ctrl-x <Up>` moves up, and `Ctrl-x <Down>` moves down. If key doesn't correspond to a direction, the command asks for a direction key.

Summary:	Alt-⟨End⟩, Ctrl-x n	next-window
	Alt-⟨Home⟩, Ctrl-x p	previous-window
	Ctrl-x ⟨Up⟩, Ctrl-x ⟨Down⟩	move-to-window
	Ctrl-x ⟨Left⟩, Ctrl-x ⟨Right⟩	move-to-window

Resizing Windows

The easiest way to resize Epsilon windows is to use the mouse. But Epsilon also provides various ways to do this via the keyboard.

The Ctrl-x + key runs the command `enlarge-window-interactively`. After you invoke the command, point to a window border using the arrow keys. The indicated window border moves so as to make the current window larger. You can keep pressing arrow keys to enlarge the window. To switch from enlarging to shrinking, press the minus key. The command Ctrl-x - works like Ctrl-x +, but starts out shrinking instead of enlarging. Whenever the window looks the right size, press ⟨Enter⟩ to leave the command.

You can use several other Epsilon commands to resize windows. The Ctrl-⟨PgUp⟩ key enlarges the current window vertically, and the Ctrl-⟨PgDn⟩ key shrinks the current window vertically. They do this by moving the mode line of the window above them up or down, if possible. Otherwise, the current window's mode line moves up or down, as appropriate.

You can also enlarge and shrink windows horizontally. The `enlarge-window-horizontally` command on Ctrl-x @ enlarges the current window by one column horizontally and the `shrink-window-horizontally` command shrinks it. They do this by moving the left boundary of the current window left or right, if possible. Otherwise, the current window's right boundary moves, as appropriate. You can use a numeric prefix with these commands to adjust by more than one line or column, or in the opposite direction.

Summary:	Ctrl-x +	enlarge-window-interactively
	Ctrl-x -	shrink-window-interactively
	Ctrl-⟨PgUp⟩, Ctrl-x ^	enlarge-window
	Ctrl-⟨PgDn⟩	shrink-window
	Ctrl-x @	enlarge-window-horizontally
		shrink-window-horizontally

4.7.4 Customizing the Screen

Epsilon displays tabs in a file by moving over to the next tab stop column. Epsilon normally spaces tabs every four or eight columns, depending on the mode. You can change the tab stop spacing by setting the variable `tab-size`. Another method is to use the `set-tab-size` command, but this can only set the tab size in the current buffer. To change the default value for new buffers, set the variable using the `set-variable` command.

Many indenting commands take the tab size into account when they indent using spaces and tabs. See page 79 for information on the indenting commands.

Epsilon can display special characters in four ways. Epsilon normally displays control characters with a ^ prefix indicating a control character (except for the few control characters like ^I that have a special meaning—^I, for example, means ⟨Tab⟩). It displays other characters with their graphic symbol.

In mode 0, Epsilon displays Meta characters (characters with codes between 160 and 255) by prefixing to them a “M-”, e.g., Meta C appears as “M-C”. Epsilon display Control-meta characters (codes 128 to 159)

by prefixing to them “M-^”, e.g., “M-^C”. Epsilon displays most control characters (codes 0 to 31) by prefixing to them a caret, e.g., Control C appears as “^C”. Characters with codes from 32 to 127 or above 255 are displayed as-is.

In mode 1, Epsilon displays graphic symbols for all the above characters, instead of using a prefix as in ^A (except for the few that have a special meaning, like `<Tab>` or `<Newline>`). Characters with codes from 32 to 127 or above 255 are displayed as-is.

In mode 2, Epsilon displays control characters and all characters with codes above 127 by their hexadecimal values, with an “x” before them to indicate hex.

In mode 3, which is the default, Epsilon displays control characters as “^C”, and uses a graphic symbol for other characters, as described above.

The `set-show-graphic` command on `Ctrl-F6` cycles among these four modes of representation. Providing a numeric argument of 0, 1, 2, or 3 selects the corresponding mode.

The command `change-show-spaces` on `Shift-F6` makes spaces, tabs, newline, and various Unicode space-like characters (including zero-width spaces) in the buffer visible, by using special graphic characters for each. Pressing it again makes these characters invisible. The command sets the buffer-specific variable `show-spaces`.

Set the buffer-specific variable `draw-line-numbers` to 1 if you want Epsilon to display line numbers. Each line’s number will appear to its left, in a field whose width is specified by the `line-number-width` variable. See the description of `draw-line-numbers` for details on its line number formatting options. (For line numbers in printed output, see the `print-line-numbers` variable.)

Epsilon will usually display a message in the echo area for at least one second before replacing it with a new message. You can set this time with the `see-delay` variable. It contains the number of hundredths of a second that a message must remain visible, before a subsequent message can overwrite it. Whenever you press a key with messages pending, Epsilon skips right to the last message and puts that up. (Epsilon doesn’t stop working just because it can’t put up a message; it just remembers to put the message up later.)

Epsilon for Windows can draw a rectangle around the current line to increase its visibility and make it easier to find the cursor. Set the `draw-focus-rectangle` variable nonzero to enable this. Set the `draw-column-markers` variable if you want Epsilon for Windows to draw a vertical line at a particular column or columns, to make it easier to edit text that must be restricted to certain columns. (Also see auto-fill mode described on page 77.)

The `set-display-characters` command lets you alter the various characters that Epsilon uses to construct its display. These include the characters Epsilon uses in some of the display modes set by `set-show-graphic`, the line-drawing characters that form window borders, and the characters it uses to construct the scroll bar. (Epsilon for Windows only uses the first of these groups.) The command displays a matrix of possible characters, and guides you through the selection process.

Cursor Shapes

You can set variables to modify the text cursor shape (the caret) that Epsilon displays in different situations.

Epsilon for Windows gets the cursor shape from one of four variables, depending upon whether or not Epsilon is in overwrite mode, and whether or not the cursor is positioned in virtual space. (See the description of the `virtual-space` variable on page 42.)

Variable	In overwrite mode?	In virtual space?
<code>normal-gui-cursor</code>	No	No
<code>overwrite-gui-cursor</code>	Yes	No
<code>virtual-insert-gui-cursor</code>	No	Yes
<code>virtual-overwrite-gui-cursor</code>	Yes	Yes

Epsilon for X11 also uses these settings, but only if you've provided an `Epsilon.cursorstyle:1` resource to let it set the cursor shape (see page 6).

Each variable contains a code that specifies the height and width of the caret, as well as a vertical offset, each expressed as a percentage of the character dimensions. Values close to 0 or 100 are absolute pixel counts, so a width of 98 is two pixels smaller than a character. A width of exactly zero means use the default width.

All measurements are from the top left corner of the character. A nonzero vertical offset moves the caret down from its usual starting point at the top left corner.

In EEL programs, you can use the `GUI_CURSOR_SHAPE()` macro to combine the three values into the appropriate code; it simply multiplies the height by 1000 and the offset by 1,000,000, and adds both to the width. So the default Windows caret shape of `GUI_CURSOR_SHAPE(100, 2, 0)`, which specifies a height of 100% of the character size and a width of 2 pixels, is encoded as the value 100,002. The value 100100 provides a block cursor, while 99,002,100 makes a good underline cursor. (It specifies a width of 100%, a height of 2 pixels, and an offset of 99 putting the caret down near the bottom of the character cell.) The `CURSOR_SHAPE()` macro serves a similar purpose for text mode versions of Epsilon.

Epsilon for Windows Console gets its cursor shapes from four other variables. As with the GUI variables, it selects one based upon whether or not Epsilon is in overwrite mode, and whether or not the cursor is positioned in virtual space.

Variable	In overwrite mode?	In virtual space?
<code>normal-cursor</code>	No	No
<code>overwrite-cursor</code>	Yes	No
<code>virtual-insert-cursor</code>	No	Yes
<code>virtual-overwrite-cursor</code>	Yes	Yes

Each of these variables contains a code that specifies the top and bottom edges of the cursor, such as 3006, which specifies a cursor that begins on scan line 3 and extends to scan line 6 on a character box. The topmost scan line is scan line 0.

Scan lines above 50 in a cursor shape code are interpreted differently. A scan line number of 99 indicates the highest-numbered valid scan line (just below the character), 98 indicates the line above that, and so forth. For example, a cursor shape like 1098 produces a cursor that extends from scan line 1 to the next-to-last scan line, one scan line smaller at top and bottom than a full block cursor.

Summary:	Ctrl-F6	<code>set-show-graphic</code>
	Shift-F6	<code>change-show-spaces</code>
		<code>set-tab-size</code>
		<code>set-display-characters</code>

4.7.5 Fonts

The `set-font` command changes the font Epsilon uses, by displaying a font dialog box and letting you pick a new font. Modifying the `font-fixed` variable is another way to set the font.

Epsilon also applies the styles bold, italic and underlined to the selected font for certain types of text, such as comments. Epsilon treats these styles as if they were part of the foreground color of a particular type of text. The `set-color` command lets you set which types of text receive which styles. Also see the variables `font-styles` and `font-styles-tolerance`.

Epsilon for Unix supports setting the font under X11, using `set-font` or `font-fixed`, but not the remaining commands and settings in this section.

You can specify a specific font for use in printing with the `set-printer-font` command. Similarly, the `set-dialog-font` command lets you specify what font to use for Epsilon's dialog windows (like the one `bufed` displays). There are also corresponding variables `font-printer` and `font-dialog`.

The command `change-font-size` supplements `set-font` by providing additional font choices. Some Windows fonts include a variety of character cell widths for a given character cell height. Commands like `set-font` utilize the standard Windows font dialog, which doesn't provide any way to select these alternate widths. The `change-font-size` command lets you choose these fonts.

The `change-font-size` command doesn't change the font name, or toggle bold or italic. You'll need to use the `set-font` command to do that.

Instead, `change-font-size` lets you adjust the height and width of the current font using the arrow keys. You can abort to restore the old font settings, or press `(Enter)` or `(Space)` to keep them. This is a handy way to shrink or expand the font size. A width or height of 0 means use a suitable default.

Summary:

```
set-font
set-printer-font
set-dialog-font
change-font-size
```

4.7.6 Setting Colors

This section describes how to set colors in Epsilon. Epsilon comes with many built-in color schemes. Each *color scheme* tells Epsilon what color to use for each *color class*. Color classes correspond to the different parts of the screen. There are separate color classes for normal text, highlighted text, text in the echo area, syntax-highlighted comments, and so forth. (See below for a partial list.)

Use the `set-color` command to select a color scheme from the list of available color schemes. You can also customize a color scheme by selecting one, selecting a color class within it, and then using the buttons to select a different foreground or background color, or toggle bold, italic, or underlined styles. The available styles depend on the selected font, as controlled by the `font-styles-tolerance` variable.

You can press `+` and `-` to expand or collapse categories in the tree of color classes. In dialog-based versions of `set-color`, the `(Right)` and `(Left)` keys also expand and collapse categories. In most versions, you can also press `Ctrl-s` or `Ctrl-r` to search for a color class by name.

Epsilon remembers the name of one color scheme for use on text mode displays with only 8 or 16 possible color choices, and a separate scheme for environments like Windows or X11 where it can display all possible colors. (It also maintains separate schemes for monochrome displays, and for when Epsilon runs as a Unix terminal program within an `xterm` and the `USE_DEFAULT_COLORS` environment variable is defined; the latter enables a special color scheme that's designed to inherit the background and foreground colors of the underlying `xterm`.)

When you've turned off window borders with the `toggle-borders` command, Epsilon uses color schemes with particular, fixed names. See page 115.

Another method of customizing a color scheme is to create an EEL file like `stdcolor.e`. The file `stdcolor.e` defines all Epsilon's built-in color schemes. You can use one of these as a model for your own color scheme. See page 425 for the syntax of color scheme definitions. You can use the `export-colors` command to build an EEL file named `mycolors.e` that contains all Epsilon's current color definitions for the current color scheme. (With a numeric argument, it lists all schemes.)

The Windows Console and Unix terminal versions of Epsilon are limited to the sixteen standard colors for foreground and background, for a total of 256 possible color combinations, while the Windows GUI and X11 versions have no such limitation. Internally, all versions of Epsilon store 32 bits of color information for the foreground and background of each color class. The console and terminal versions convert back to 4 bits of foreground and background when displaying text. In these environments, there are no buttons for selecting a foreground or background color. Instead, the arrow keys select colors.

The `set-color` command displays a short description of each color class as you select it. Here we describe a few of the color classes in more detail:

`text` Epsilon puts the text of an ordinary buffer in this color. But if Epsilon is doing code coloring in a buffer, it uses the color classes defined for code coloring instead. For instance, C++ and Java files both use C mode, and the color classes defined for C mode all start with “c-” and appear farther down in the list of color classes.

`mode-line` Epsilon uses this color for the text in the mode line of a tiled window.

`horiz-border` Epsilon uses this color for the line part of the mode line of a tiled window.

`vert-border` Epsilon uses this color for the vertical border it draws between tiled windows.

`after-exiting` Some console versions of Epsilon try to leave the screen in this color when you exit. Epsilon normally sets this color when it starts up, based on the screen’s colors before you started Epsilon. Set the `restore-color-on-exit` variable to zero to disable this behavior, so you can set the color explicitly and preserve the change in your state file.

`debug-text` The EEL debugger uses this color when it displays EEL source code.

`default` Epsilon initializes any newly-defined color classes (see page 419) with this color.

`screen-border` Epsilon sets the border area around the screen or window to match this color’s background. Epsilon only uses the background part of this color; the foreground part doesn’t matter.

`screen-decoration` Epsilon for Windows can draw a focus rectangle or column markers. The foreground color specified here determines their color. See the `draw-focus-rectangle` and `draw-column-markers` variables.

`pull-highlight` The `pull-word` command uses this color for its highlighting.

Summary:	<code>set-color</code>
	<code>export-colors</code>

4.7.7 Code Coloring

Epsilon does syntax-based highlighting for many different programming languages. Set the buffer-specific variable `want-code-coloring` to 0 to disable this feature or run the `change-code-coloring` command. To change the colors Epsilon uses, see the previous section. (Because certain modes like Perl and HTML use coloring to quickly parse language syntax, if you don’t want to see the coloring it’s often better to change the color selections so they’re identical instead of disabling code coloring entirely.)

If you use a very old and slow computer, you may need to tell Epsilon to do less code coloring, in order to get acceptable response time. Set the variable `minimal-coloring` to 1 to tell Epsilon to look only for comments, preprocessor lines, strings, and character constants when coloring. Epsilon will color all

identifiers, functions, keywords, numbers and punctuation the same, using the `c-ident` color class for all. This makes code coloring much faster.

When Epsilon begins coloring in the middle of a buffer, it has to determine whether it's inside a comment by searching back for comment characters. If you edit extremely large C files with few block comments, you can speed up Epsilon by telling it not to search so far. Set the variable `color-look-back` to the number of characters Epsilon should search through before giving up. Any block comments larger than this value may not be colored correctly. A value of zero (the default) lets Epsilon search as far as it needs to, and correctly colors comments of any size.

When Epsilon isn't busy acting on your keystrokes, it looks through the current buffer and assigns colors to the individual regions of text, so that Epsilon responds faster as you scroll through the buffer. For smoother performance, Epsilon doesn't begin to do this until it's been idle for a certain period of time, contained in the `idle-coloring-delay` variable. This holds the number of hundredths of a second to wait before computing more coloring information. By default, it's 100, so Epsilon waits one second. Set it to -1 to disable background code coloring.

Normally Epsilon colors buffers as needed. You can set Epsilon to instead color the entire buffer the first time it's displayed. Set the variable `color-whole-buffer` to the size of the largest buffer you want Epsilon to entirely color at once.

Summary:

`change-code-coloring`

4.7.8 Window Borders

Under Windows and X11, you can control the title of Epsilon's main window. The variables `window-caption-file`, `window-caption-buffer`, and `window-caption` control what appears in Epsilon's title bar.

Use the command `set-display-look` to make Epsilon's window decoration and screen appearance resemble that of other editors. It displays a menu of choices. You can select Epsilon's original look, Brief's look, the look of the DOS Edit program (the same as the QBasic program), or the look of the Borland IDE.

The command `toggle-borders` removes the lines separating Epsilon's windows from one another, or restores them.

When there are no window borders, Epsilon provides each window with its own separate color scheme, in place of the single one selected by `set-color`. (You can still use `set-color` to set the individual colors in a color scheme, but Epsilon doesn't care which particular color scheme you select when it displays the contents of individual windows. It does use your selected color scheme for other parts of the screen like the echo area or screen border.)

The color schemes Epsilon uses for borderless windows have names like `"window-black"`, `"window-blue"` and so forth. Epsilon assigns them to windows in order. You can remove one from consideration using the `delete-name` command, or create a new one using EEL (see page 425).

The rest of this section describes some of the variables set by the above commands. The `set-display-look` command in particular does its work entirely by setting variables. You can make Epsilon use a custom display look by setting these variables yourself. The variables also allow some customizations not available through the above commands.

The `echo-line` variable contains the number of the screen line on which to display the echo area. The `avoid-top-lines` and `avoid-bottom-lines` variables tell Epsilon how many screen lines at the top and bottom of the screen are reserved, and may not contain tiled windows. By default, `echo_line` contains the

number of the last screen line, `avoid-top-lines` is zero, and `avoid-bottom-lines` is one, to make room for the echo area.

To Epsilon display text in the echo area whenever it's idle, set the variables `show-when-idle` and `show-when-idle-column`. See their online documentation for details.

To position the echo area at the top of the screen, set `echo-line` and `avoid-bottom-lines` to zero and `avoid-top-lines` to one. (If you're using a permanent mouse menu, set `echo-line` and `avoid-top-lines` one higher.)

To completely fill the screen with text, toggle borders off and set `avoid-bottom-lines` and `avoid-top-lines` to zero. Whenever Epsilon needs to display text in the echo area, it will temporarily overwrite the last screen line for a moment, and then return to showing buffer text on every line.

You can customize the position and contents of the mode line Epsilon displays for ordinary tiled windows by setting variables. These variables all start with "mode-". See the online help for `mode-format` for details. Also see the `full-path-on-mode-line` variable.

You can set several variables to put borders around the screen. If you want Epsilon to always display a window border at the right edge of the screen, set the variable `border-right` nonzero. (The `toggle-scroll-bar` command, which turns on permanent scroll bars for all windows, sets this variable.) Epsilon displays a border at the left screen edge if `border-left` has a nonzero value. Similarly, `border-top` and `border-bottom` variables control borders at the top and bottom edges of the screen, but only if a tiled window reaches all the way to that edge of the screen. (A menu bar might be in the way.) All these variables are zero by default. (Toggling all window borders off with the `toggle-borders` command overrides these variables.) If the `border-inside` variable is nonzero (as it is by default), Epsilon displays a border between side-by-side windows. Set it to zero to eliminate these borders. (The `toggle-borders` command sets this variable, among other things.)

Summary:

`set-display-look`

4.7.9 The Bell

Sometimes Epsilon will ring the computer's bell to alert you to certain conditions. (Well, actually it sounds more like a beep, but we call it a bell anyway.) You can enable or disable the bell completely by setting the `want-bell` variable. Epsilon will never try to beep if `want-bell` has a value of zero.

For finer control of just when Epsilon rings the bell, you can set the variables listed in figure 4.6 using the `set-variable` command, described on page 159. A nonzero value means Epsilon will ring the bell when the indicated condition occurs. By default, all these variables but `bell-on-abort` have the value 1, so Epsilon rings the bell on almost all of these occasions.

In some environments, the `beep-duration` variable specifies the duration of the beep, in hundredths of a second. The `beep-frequency` variable specifies the frequency of the bell in hertz.

Instead of making a sound for the bell, you can have Epsilon invert the mode line of each window for a time according to the value of `beep-duration` by setting `beep-frequency` to zero, and `beep-duration` to any nonzero value.

Under Windows, Epsilon doesn't use the `beep-duration` or `beep-frequency` variables. It uses a standard system sound instead. Under Unix, Epsilon recognizes a `beep-frequency` of zero and flashes the screen in some fashion, but otherwise ignores these variables.

Variable	When Epsilon Beeps, if Nonzero
<code>bell-on-abort</code>	You abort a command, press an unbound key, or Epsilon reports an error.
<code>bell-on-autosave-error</code>	Autosaving can't write files.
<code>bell-on-bad-key</code>	You press an illegal option at a prompt.
<code>bell-on-completion</code>	Completion finds no matches.
<code>bell-on-date-warning</code>	Epsilon notices that a file has changed on disk.
<code>bell-on-read-error</code>	Epsilon cannot read a file.
<code>bell-on-search</code>	Search finds no more matches.
<code>bell-on-write-error</code>	Epsilon cannot write a file.

Figure 4.6: Variables that control when Epsilon rings the bell

4.8 Buffers and Files

4.8.1 Buffers

The `Ctrl-x b` command prompts you for a buffer name. The command creates a buffer if one with that name doesn't already exist, and connects the buffer to the current window.

The `new-file` command creates a new buffer and marks it so that Epsilon will prompt for its file name when you try to save it. It doesn't prompt for a buffer name, unlike `Ctrl-x b`, but chooses an unused name.

You can customize the behavior of the `new-file` command by setting the variables `new-file-mode` and `new-file-ext`. The `new-file-mode` variable contains the name of the mode-setting command Epsilon should use to initialize new buffers; the default is the `c-mode` command. The `new-file-ext` variable contains the extension of the file name Epsilon constructs for the new buffer; its default is `".c"`.

To get a list of the buffers, type `Ctrl-x Ctrl-b`. This runs the `bufed` (for buffer edit) command, described fully on page 143. Basically, `bufed` lists your buffers, along with their sizes and the files (if any) contained in those buffers. You can then easily switch to any buffer by positioning point on the line describing the buffer and pressing the `<Space>` key. The `bufed` command initially positions point on the buffer from which you invoked `bufed`. Press `Ctrl-g` if you decide not to switch buffers after all.

The `Ctrl-x k` command eliminates a buffer. It asks you for a buffer name and gets rid of it. If the buffer has unsaved changes, the command warns you first.

The `Ctrl-x Ctrl-k` command eliminates the current buffer, just like `Ctrl-x k`, but without asking which buffer you want to get rid of. The `kill-all-buffers` command discards all user buffers.

Whenever Epsilon asks you for a buffer name, it can do completion on buffer names, and will list matches in a pop-up window if you press `'?`.

Another way to switch buffers is to press `Ctrl-Tab`. This command switches to the buffer you last used. If you press `<Tab>` again while still holding down `Ctrl`, you can switch to still older buffers. Hold down `Shift` as well as `Ctrl` to move in the reverse order. You can press `Ctrl-g` to abort and return to the original buffer.

You can also change to another buffer using the `next-buffer` and `previous-buffer` commands. They select the next (or previous) buffer and connect it to the current window. You can cycle through all the buffers by repeating these commands. You can type `F12` and `F11`, respectively, to run these commands. If your keyboard doesn't have these keys, you can also type `Ctrl-x >` and `Ctrl-x <`.

Summary:	<code>Ctrl-x b</code>	<code>select-buffer</code>
	<code>Ctrl-x Ctrl-b</code>	<code>bufed</code>

Ctrl-x k	kill-buffer
Ctrl-x Ctrl-k	kill-current-buffer
Ctrl-⟨Tab⟩	switch-buffers
F12, Ctrl-x >	next-buffer
F11, Ctrl-x <	previous-buffer
	kill-all-buffers
	new-file

4.8.2 Files

Reading Files

The Ctrl-x Ctrl-f key runs the `find-file` command. It prompts you for a file name. First, it scans the current buffers to see if any of them contain that file. If so, the command connects that buffer to the current window. Otherwise, the command creates a buffer with the same name as the file, possibly modified to make it different from the names of existing non-empty buffers, then reads the file into that buffer. Most people consider `find-file` the command they typically use to edit a new file, or to return to a file read in previously.

Normally Epsilon examines the file's contents to determine if it's a binary file, or in Unix or Macintosh format. If you prefix a numeric argument to `find-file`, Epsilon asks you for the correct format, as described on page 122, and for the name of an encoding. (Use the `auto-detect` encoding if you only want to force binary, Unix, Macintosh, or Windows format.)

The `find-file` command examines a file's name and contents to determine an appropriate language mode for it. For instance, files with a `.c` extension are put in C mode. You can override this decision with a "file variable". See page 127. You can use the `reset-mode` command at any time to make Epsilon repeat that process, setting the buffer to a different mode if appropriate. It can be handy after you've temporarily switched to a different mode for any reason, or after you've started creating a new file with no extension and have now typed the first few lines, enough for Epsilon to auto-detect the proper mode.

If you type ⟨Enter⟩ without typing any file name when `find-file` asks for a file, it runs `dired` on the current directory. If you give `find-file` a file name with wild card characters, or a directory name, it runs the `dired` command giving it that pattern. See page 139 for a description of the very useful `dired` command. Also see page 125 for information on related topics like how to type a file name with ⟨Space⟩ characters, customize the way Epsilon prompts for files, and so forth.

By default, at most prompts for file names like `find-file`'s, Epsilon types in for you the directory portion of the current file. For example, suppose the current buffer contains a file named `"\src\new\ll.c"`. If you invoke `find-file`, Epsilon will type in `"\src\new\"` for you. This comes in handy when you want to read another file in the same directory as the current file. You can simply begin typing another file name if you want Epsilon to ignore the pre-typed directory name. As soon as Epsilon notices you're typing an absolute file pathname, it will erase the pre-typed directory name. See page 125 for details.

You can change the current directory with the `cd` command on F7. It prompts for a new current directory, and then displays the full pathname of the selected current directory. You can type the name of a new directory, or just type ⟨Enter⟩ to stay in the current directory. When you supply a file name, Epsilon interprets it with respect to the current directory unless it begins with a slash or backslash. If you specify a drive as part of the directory name, Epsilon will set the current drive to the indicated drive, then switch to the indicated directory. Press Alt-e when prompted for a directory name, and Epsilon will insert the name of the directory containing the current file. (If you start a concurrent process, Epsilon can link its current directory to Epsilon's; see the `use-process-current-directory` variable for details.)

The `insert-file` command on `Ctrl-x i` prompts for the name of a file and inserts it before point. It sets the mark before the inserted text, so you can kill it with `Ctrl-w`. (Also see the `insert-file-remembers-file` variable.)

The `find-linked-file` command on `Ctrl-x Ctrl-l` looks for a file name in the current buffer, then finds that file. It works with plain text files, and also understands `#include` in C-like buffers, `` in HTML-like buffers, and various other mode-specific conventions. It also understands file names that start with an environment variable, such as `%HOMEPATH%\Documents` under Windows, or `$PERL5LIB/CPAN` or `${PERL5LIB}CPAN` on macOS, Linux and FreeBSD. (See the `file-interpret-env-vars` variable to disable this.)

You can highlight a file name first whenever `find-linked-file`'s automatic parsing of file names isn't right. In a process buffer, it looks for error messages, not file names (unless you've first highlighted a file name), and sets the current error message (as used by `next-error`) to the current line.

Epsilon uses a built-in list of directories to search for `#include` files; you can set the `include-directories` variable to add to that list. (The `copy-include-file-name` command also uses this list of directories.) For files with a `.lst` extension, it assumes the current line holds a file name, instead of searching for a pattern that matches a typical file name. This is one way to more easily manage files in a project that are in many different directories.

The key `Ctrl-x Ctrl-v` runs the `visit-file` command. It prompts you for a file name. If the file exists, the command reads it into the current buffer, and positions point at the beginning. The command discards the old contents of the buffer, but asks before discarding an unsaved buffer. If no file with the given name exists, the command clears the current buffer. If you prefix this command with a numeric argument, the command discards the old buffer content without warning. So if you want to revert to the copy of the file on disk, disregarding the changes you've made since you last saved the buffer, press `Ctrl-u Ctrl-x Ctrl-v`, followed by `(Enter)`. Most people use this command only to explicitly manipulate the file associated with a particular buffer. To read in a file, use the `find-file` command, described above.

The `revert-file` command rereads the current file from disk. If you've made any unsaved changes, it prompts first.

If a file has an extension `.gz` or `.bz2`, indicating a compressed file, Epsilon automatically decompresses it when you read it. See the `uncompress-files` variable. This feature runs the standard utility programs `gzip` (for `.gz`) and `bzip2` (for `.bz2`); you'll need to install them if they're not already installed. For Windows, the programs provided by the Cygwin environment will work fine.

Summary:	<code>Ctrl-x Ctrl-f</code>	<code>find-file</code>
	<code>F7</code>	<code>cd</code>
	<code>Ctrl-x i</code>	<code>insert-file</code>
	<code>Ctrl-x Ctrl-v</code>	<code>visit-file</code>
		<code>revert-file</code>

Read-Only Files

Whenever you read a read-only file into a buffer using `find-file` or `visit-file`, Epsilon makes the buffer read-only, and indicates this by displaying "RO" in the modeline. Epsilon keeps you from modifying a read-only buffer. Attempts to do so result in an error message. In a read-only buffer you can use the `(Space)` and `(Backspace)` keys to page forward and back more conveniently; see the `readonly-pages` variable to disable this.

If you want to modify the buffer, you can change its read-only status with the `change-read-only` command on `Ctrl-x Ctrl-q`. With no numeric argument, it toggles the read-only status. With a non-zero

Set the buffer-specific variable `want-warn` to 0 if you don't want Epsilon to ever check the file date or warn you. Or under Windows, set the `file-date-skip-drives` variable to make Epsilon ignore file dates on specific types of drives, such as network drives or floppy disks. Epsilon never checks file dates for URLs.

You can have Epsilon remove any spaces or tabs at the end of each line, before saving a file. See the `c-delete-trailing-spaces` and `default-delete-trailing-spaces` variables.

Similarly, you can have Epsilon make sure files you save end with a line termination character like a newline by setting the `default-add-final-newline` variable. To get this behavior only for specific modes, create a variable with a name like `html-add-final-newline` and Epsilon will use its setting instead for buffers in that mode.

Epsilon automatically marks a buffer as “modified” when you change it, and shows this with a star ‘*’ at the end of the buffer’s mode line. When Epsilon writes a buffer to disk or reads a file into a buffer, it marks the buffer as “unmodified”. When you try to exit Epsilon, it will issue a warning if any buffer contains a file with unsaved changes.

You may occasionally want to change a buffer’s modified status. You can do this with the `change-modified` command. Each time you invoke this command, the modified status of the current buffer toggles, unless you invoke it with a numeric argument. A nonzero numeric argument sets the modified status; a numeric argument of zero clears the modified status.

The `save-all-buffers` command, bound to `Ctrl-x s`, goes to each buffer with unsaved changes (those marked modified), and if it contains a file, writes the buffer out to that file. See the `save-all-without-asking` variable to alter what Epsilon does when there’s an error saving a file.

The `write-region` command on `Ctrl-x w` takes the text between point and mark, and writes it to the file whose name you provide.

Summary:	<code>Ctrl-x Ctrl-s</code>	<code>save-file</code>
	<code>Ctrl-x Ctrl-w</code>	<code>write-file</code>
	<code>Alt-~</code>	<code>change-modified</code>
	<code>Ctrl-x s</code>	<code>save-all-buffers</code>
	<code>Ctrl-x w</code>	<code>write-region</code>
		<code>set-file-name</code>

Backup Files

Epsilon doesn't normally keep the previous version of a file around when you save a modified version. If you want backups of saved files, you can set the buffer-specific variable `want-backups` to 1, using the `set-variable` command described on page 159. If this variable is 1, the first time you save a file in a session, Epsilon will first preserve the old version by renaming any existing file with that name to a file with the extension “.bak”. For instance, saving a new version of the file `text.c` preserves the old version in `text.bak`. (If you delete a file’s buffer and later read the file again, Epsilon treats this as a new session and makes a new backup copy the next time you save.) If `want-backups` variable is 2, Epsilon will do this each time you save the file, not just the first time. The `backup-by-renaming` variable controls whether Epsilon backs up files by renaming them (faster) or copying them (necessary in some environments to preserve attached attributes).

You can change the name Epsilon uses for a backup file by setting the variable `backup-name`, which holds a file name template (see the next section). The default setting `%p%b.bak` uses the same path and base file name as the original file but replaces the extension with `.bak`.

Epsilon automatically saves a copy of each modified file periodically, into a separate file with a name like `#file.c.asv#`, then deletes the autosaved file when you tell it to save the original (and when Epsilon

exits). The variable `want-auto-save` controls this. Epsilon uses a template (see above) to construct the name of each auto save file, stored in the variable `auto-save-name`. Other bits in the `want-auto-save` variable let you make auto-saving more verbose, or tell Epsilon not to automatically delete auto-saved files when exiting, or when the file is saved normally.

You can alter the number of keystrokes it waits between autosaves by setting the variable `auto-save-count`. Epsilon also auto-saves after you've been idle for 30 seconds; set the `auto-save-idle-seconds` variable to alter this number. Very large buffers will never be auto-saved; see the `auto-save-biggest-file` variable to alter this.

Sometimes you may want to explicitly write the buffer out to a file for backup purposes, but may not want to change the name of the file associated with the buffer. For that, use the `copy-to-file` command on Ctrl-F7. It asks you for the name of a file, and writes the buffer out to that file, but subsequent Ctrl-x Ctrl-s's will save to the original file.

Summary: Ctrl-F7 `copy-to-file`

File Name Templates

Epsilon uses file name templates to construct one file name from another, such as constructing the name of a backup file from the original file name. Epsilon also uses these to construct command lines, for example in the various `compile-?-cmd` variables that the `compile-buffer` command uses, or with the `%` sequence in .mnu files.

Epsilon copies the template text, substituting pieces of the original file name when it encounters codes in the template, according to figure 4.8. The sequence `%r` substitutes a relative pathname to the original file name, if the file is within the current directory or its subdirectories, or an absolute pathname otherwise.

The sequence `%h` substitutes the host name, if the original file name is in URL format, otherwise nothing. The sequence `%c` substitutes the name of Epsilon's customization directory (see page 13).

The sequence `%x` substitutes the full pathname of the directory containing the Epsilon executable. The sequence `%X` substitutes the same full pathname, but this time after converting all Windows long file names making up the path to their equivalent short name aliases, if any. For example, if the Epsilon executable was in the directory `c:\Program Files\Eps14\bin\`, `%x` would use exactly that pathname, while `%X` might yield `c:\Progra~1\Eps14\bin\`. Under Unix, `%X` is the same as `%x`. Either always ends with a path separator character like `/` or `\`.

If any other character follows `%`, Epsilon puts that character into the resulting file name. You can use this, for example, to include an actual `%` character in the result by putting `%%` in the template.

Line Translation

Most Windows programs use files with lines separated by the pair of characters Return, Newline (or Control-M, Control-J), just as DOS did. But internally Epsilon separates lines with just the newline character, Ctrl-j. Epsilon normally translates between the two systems automatically when reading or writing text files in this format. When it reads a file, it removes all Ctrl-m characters, and when it writes a file, it adds a Ctrl-m character before each Ctrl-j.

Epsilon will automatically select one of several other translation types when appropriate, based on the contents of the file you edit. It automatically determines whether you're editing a regular file, a binary file, a Unix file, or an old Mac file (since OS X, Macs have used Unix-style text files), and uses the proper translation scheme. You can explicitly override this if necessary. Epsilon determines the file type by looking

		Example 1	Example 2
Code	Part	c:\dos\read.me	/usr/bin
%p	Path	c:\dos\	/usr/
%b	Base	read	bin
%e	Extension	.me	(None)
%f	Full name	c:\dos\read.me	/usr/bin
%r	Relative path (assuming current directory is	dos\read.me	/usr/bin
		c:\	/usr/mark)
%h	Host name of URL		
%c	Customization dir	...\Lugaru\Epsilon	/home/mark/.epsilon
%x	Executable path	c:\Program Files\Eps14\bin\	/opt/epsilon14.00/bin/
%X	Alias to path	c:\Progra~1\Eps14\bin\	/opt/epsilon14.00/bin/

Figure 4.8: File name template characters.

at the first few hundred thousand bytes of the file, and applying heuristics. This is quite reliable in practice. However, Epsilon may occasionally guess incorrectly. You can tell Epsilon exactly which translation scheme to use by providing a numeric argument to a file reading command like `find-file`, or a file-writing command like `save-file` or `write-file`. Epsilon will then prompt for which translation scheme to use.

The `set-line-translate` command sets this behavior for the current buffer. It prompts for the desired type of translation, and makes future file reads and writes use that translation. Epsilon will display “Binary”, “Unix”, “DOS”, or “Mac” in the mode line to indicate any special translation in effect. (It omits this when the “usual” translation is in effect: Unix files in Epsilon for Unix, DOS files in other versions.)

Set the `default-translation-type` variable if you want to force Epsilon to always use a particular type of translation when reading existing files, rather than examining their contents and choosing a suitable type. A value of 0 forces binary, 1 forces DOS/Windows, 2 forces Unix, and 3 forces Macintosh. A value of 5, the default, lets Epsilon autodetect the file type.

Set the `new-buffer-translation-type` variable if you want Epsilon to create new buffers and files with a translation type other than the default. For file names that start with `ftp://`, the `ftp-ascii-transfers` variable can change the meaning of some translation types; see its online help.

For file names in the form of a URL, Epsilon uses the `force-remote-translation-type` variable instead of `default-translation-type`. When it’s not set to 5 to request auto-detection, it makes Epsilon use one specific translation type for all remote files, bypassing auto-detection.

Setting the `fallback-remote-translation-type` variable instead lets auto-detection proceed, but sets the translation type Epsilon uses whenever it can’t determine a type by examining the file, and for new files. The default value, 5, makes Epsilon for Unix pick Unix, and Epsilon for Windows pick DOS/Windows. This variable is the remote-file equivalent of `new-buffer-translation-type`.

Host-specific variables takes precedence over both `force-remote-translation-type` and `fallback-remote-translation-type`, letting you establish separate settings for each remote system. See these variables’ full descriptions for details.

Epsilon remembers the type of translation you want in each buffer using the buffer-specific variable `translation-type`.

You can use the “`write-line-translate`” file variable to set Epsilon so it auto-detects the translation rule when reading existing files, but forces all files into a specific mode when saving them. See page 128.

Epsilon applies the following heuristics, in order, to determine a file's type. These may change in future versions.

A file that contains null bytes is considered binary. A file that has no Ctrl-m Ctrl-j pairs is considered a Unix file if it contains Ctrl-j characters, or a Macintosh file if it contains Ctrl-m. A file containing a Ctrl-m character not followed by either Ctrl-m or Ctrl-j is considered binary. So is a file containing a Ctrl-j character not preceded by a Ctrl-m as well as some Ctrl-m Ctrl-j pairs. Any other files, or files of less than five characters, are considered to be in standard DOS/Windows format (or in Epsilon for Unix, Unix format).

Bear in mind that Epsilon makes all these decisions after examining only the first few hundred thousand bytes of a file, and phrases like "contains null bytes" really mean "contains null bytes in its first few hundred thousand characters." When Epsilon chooses a file type based on text far from the start of the file, so that the reason for the choice may not be obvious, it displays a message explaining why it picked that translation type. The `file-read-kibitz` variable controls this.

Summary: set-line-translate

DOS/OEM Character Set Support

Windows programs typically use a different character set than do DOS programs, or programs that run in a Windows Console environment. The DOS character set is known as the DOS/OEM character set, and includes various line drawing characters and miscellaneous characters not in the Windows/ANSI set. The Windows/ANSI character set includes many accented characters not in the DOS/OEM character set. Epsilon for Windows uses the Windows/ANSI character set (with most fonts). Epsilon for Windows Console uses a DOS/OEM character set by default, but see the `console-ansi-font` variable.

The Unicode-based commands described on page 136 are usually a better choice than the older commands documented in this section.

The `oem-to-ansi` command converts the current buffer from the DOS/OEM character set to the Windows/ANSI character set. The `ansi-to-oem` command does the reverse. If any character in the buffer doesn't have a unique translation, these commands warn before translating, and move to the first character without a unique translation.

The `find-oem-file` command reads a file using the DOS/OEM character set, translating it into the Windows/ANSI character set, and arranges things so when you save the file, the reverse translation automatically occurs.

The commands in this section provide a subset of the functionality available with the Unicode-based commands described on page 136. The `oem-to-ansi` command is similar to the `unicode-convert-from-encoding` command. Specify an encoding such as "cp850" or "cp437", using the code page number shown by the "chcp" command at a Windows command prompt. Similarly, the `ansi-to-oem` command is like the `unicode-convert-to-encoding` command. The `find-oem-file` command is like invoking `find-file` with a numeric prefix argument, so it asks for line translation and encoding options, and specifying the DOS/OEM encoding as above. See page 136 for details on setting a default code page and similar options.

Summary: oem-to-ansi
ansi-to-oem
find-oem-file

File Name Prompts

By default, Epsilon interprets environment variables when you type file names, using the syntax `%TEMP%` under Windows, and `$TEMP` or `${TEMP}` on macOS, Linux and FreeBSD. Keys like `?` and `<Space>` complete on environment variable names just as they do on file names.

The list of names Epsilon recognizes between `%` characters also includes Windows shell folder names. Epsilon prefixes these with an `@` to prevent collisions with environment variable names.

On all platforms, you can use the name `@ecustom` as if it were an environment variable name, to refer to Epsilon's customization directory.

You can define your own names to be substituted just like environment variable names, by putting them in a text file named `subst_vars` in your Epsilon customization directory (see page 13). Each line in the file should have the format `varname=value`, where `varname` may contain letters, digits, `_`, and `@`. Your variables can name either a directory (which must end with a `\` or `/`) or a specific file. For example:

```
@recipies=/home/bob/misc/recipies/
@todo=/home/bob/Documents/todo
```

When such a variable's value is a list of directories, Epsilon converts it into a file pattern, so a pattern like `$INCLUDE/std*.h` (or, in Windows syntax, `%INCLUDE%\std*.h`) matches files in any directory on the `INCLUDE` path. (Actually Epsilon just checks if the value contains the directory separator character: semicolon for Windows and colon elsewhere.) Epsilon transforms such values using extended file pattern syntax, so any file name following matches in any of the listed directories.

Your values can also use extended file patterns directly. For example, say you frequently need to grep or perform text replacements on all `.c`, `.h`, and `.cpp` files in the `c:\myprogram` hierarchy. You could define a variable as follows, and then just type `%@src%` at those commands' file pattern prompts.

```
@src=c:\myprogram\**\*.{[ch],cpp}
```

You can disable such variable substitution by typing `"` before or within the name, or turn off some or all aspects of it by setting bits in the `file-interpret-env-vars` variable.

Epsilon caches the list of variable names and values in the `-subst-vars` buffer. Delete that buffer with the `kill-buffer` command on `Ctrl-x k` if you want Epsilon to load changes to your Windows shell folder settings or file of custom variables.

You can customize many other aspects of Epsilon's behavior when prompting for file names.

By default, many commands in the Windows version of Epsilon use the standard Windows common file dialog, but only when you invoke them from a menu or the tool bar. When you invoke these commands using their keyboard bindings, they use the same kind of dialog as other Epsilon prompts.

Set `want-common-file-dialog` to 2 if you want Epsilon to use the common file dialog whenever it can. Set `want-common-file-dialog` to 0 to prevent Epsilon from ever using this dialog. The default value of 1 produces the behavior described above. You can use the `force-common-file-dialog` command to toggle whether Epsilon uses a dialog for the next command only.

The Windows common file dialog includes a list of common file extensions. You can customize this list by editing the `file.filter.txt`, putting your own version in your customization directory (see page 13). See the comments in that file for more information. You can also customize which directory this dialog uses, and how Epsilon remembers that choice; see the `common-open-use-directory` variable.

All the remaining features and variables described in this section have no effect when Epsilon uses the standard Windows dialog; they only modify Epsilon's own file dialogs. Similarly, the standard Windows common file dialog cannot interpret environment variables and other names as described above.

The `prompt-with-buffer-directory` variable controls how Epsilon uses the current directory at file prompts. When this variable is 2, the default, Epsilon inserts the current buffer's directory at many file prompts. This makes it easy to select another file in the same directory. You can edit the directory name, or you can begin typing a new absolute pathname right after the inserted pathname. Epsilon will delete the inserted pathname when it notices your absolute pathname. This behavior is similar to Gnu Emacs's. (See the `yank-options` variable to modify how Epsilon deletes the inserted pathname.)

A setting of 3 makes Epsilon insert the current buffer's directory in the same way, but prevents Epsilon from automatically deleting the inserted pathname if you type an absolute one.

When `prompt-with-buffer-directory` is 1, Epsilon temporarily changes to the current buffer's directory while prompting for a file name, and interprets file names relative to the current directory. This behavior is similar to the "pathname.e" extension available for previous versions of Epsilon.

When `prompt-with-buffer-directory` is 0, Epsilon doesn't do anything special at file prompts. This was Epsilon's default behavior in previous versions.

The `grep` and `file-query-replace` commands use a separate variable `grep-prompt-with-buffer-directory` for their file patterns, with the same meaning as above. By default it's 1.

During file name completion, Epsilon can ignore files with certain extensions. The `ignore-file-extensions` variable contains a list of extensions to ignore. By default, this variable has the value `'|.obj|.exe|.o|.b|'`, which makes file completion ignore files that end with `.obj`, `.exe`, `.o`, and `.b`. Each extension must appear between `'|'` characters. You can augment this list using the `set-variable` command, described on page 159.

Similarly, the `only-file-extensions` variable makes completion look only for files with certain extensions. It uses the same format as `ignore-file-extensions`, a list of extensions surrounded by `|` characters. If the variable holds a null pointer, Epsilon uses `ignore-file-extensions` as above.

Completion also restricts its matches using the `ignore-file-basename` and `ignore-file-pattern` variables, which use patterns to match the names of files to be excluded. If the pattern the user types doesn't match any files, due to any of the various exclusion variables, Epsilon temporarily removes all exclusions and lists matching files again.

When Epsilon prompts for a file name, the `<Space>` key performs file name completion on what you've typed. To create a new file with spaces in its name, you must quote the space characters by typing `Ctrl-q` before each one, while entering the name, or type `"` characters around the file name (or any part containing spaces).

At any Epsilon prompt (not just file prompts), you can type `Alt-e` to retrieve your previous response to that prompt. `Alt-⟨Up⟩` or `Ctrl-Alt-p` show a list of previous responses. See page 30 for complete details. `Alt-⟨Down⟩` or `Ctrl-Alt-n` let you easily copy text from the buffer into the prompt (useful when the buffer contains a file name or URL). See page 28 for more information. At most file name prompts, `Alt-g` will retrieve the name of the current buffer's file.

When Epsilon for Windows shows a dialog containing a list of previous responses, or files matching a pattern, the list may be too wide for the dialog. You can generally resize the dialog by simply dragging its border. This works for most Epsilon dialogs. Epsilon will automatically remember the size of each dialog from session to session.

File Name Case

When retrieving file names from some file systems, Epsilon automatically translates the file names to lower case. Epsilon uses various different rules for determining when to convert retrieved file names to lower case, and when two file names that differ only by case refer to the same file.

Epsilon distinguishes between three types of file systems:

On a case-sensitive file system, MyFile, MYFILE, and myfile refer to three different files. Unix file systems are normally case-sensitive.

On a case-preserving (but not case-sensitive) file system, MyFile, MYFILE, and myfile all refer to the same file. But if you create a file as MyFile, the file system will display that file as MyFile without altering its case. VFAT, NTFS, and HFS file systems used in Windows and Mac OS are case-preserving.

On a non-case-preserving file system, MyFile, MYFILE, and myfile all refer to the same file. Moreover, the operating system converts all file names to upper case. So no matter how you create the file, the operating system always shows it as MYFILE. DOS's FAT file system is non-case-preserving. When Epsilon displays a file name from such a file system, it changes the file name to all lower case.

Epsilon for Windows asks the operating system for information on each drive, the first time the drive is accessed. Epsilon for Unix assumes all file systems are case-sensitive (for Mac OS, case-preserving), and the rest of this section does not apply.

You can tell Epsilon to use particular rules for each drive on your system by defining an environment variable. The MIXEDCASEDRIVES environment variable should contain a list of drive letters or ranges. If the variable exists and a lower case letter like k appears in it, Epsilon assumes drive K: has a Unix-style case-sensitive file system. If the variable exists and an upper case letter like J appears in it, Epsilon assumes drive J: is not case-preserving or case-sensitive, like traditional FAT drives. If the variable exists but a drive letter does not appear in it, Epsilon assumes the drive has a case-preserving but not case-sensitive file system like NTFS, HPFS, or VFAT drives.

If, for example, drives h:, i:, j:, and p: access Unix filesystems over a network, drive q: accesses a server that uses a FAT filesystem, and other drives use a VFAT filesystem (local drives under Windows, for example), you could set MIXEDCASEDRIVES to h-jpQ. When Epsilon finds a MIXEDCASEDRIVES variable, it assumes the variable contains a complete list of such drives, and doesn't examine filesystems as described. If an EPSMIXEDCASEDRIVES configuration variable exists, that overrides any MIXEDCASEDRIVES environment variable that may be found. (Note that MIXEDCASEDRIVES appears in the environment under all operating systems, while EPSMIXEDCASEDRIVES is a configuration variable must be put in the registry under Windows. See page 10 for details.)

You can set the variable `preserve-filename-case` nonzero to tell Epsilon to use the case of filenames exactly as retrieved from the operating system. By default, Epsilon for Windows changes all-uppercase file names to lower case, except on case-sensitive file systems. The variable also controls case conversion rules when Epsilon picks a buffer name for a file, and related settings.

4.8.3 File Variables

The find-file command examines a file's name and contents to determine an appropriate language mode for it. For instance, files with a .c extension are put in C mode. You can override this decision with a "file variable".

These are specially-formatted lines at the top or bottom of a file that indicate the file should use a particular language mode, tab size, or other setting. For example, you can put `-- mode: VBasic --` anywhere on the first line of a file to force Epsilon to Visual Basic mode, or write `-- tab-size: 3 --` to make Epsilon use that tab size setting.

Epsilon recognizes a syntax for file variables that's designed to be generally compatible with Emacs. The recognized formats are as follows. First, the first line of the file (or the second, if the first starts with #!, to accommodate the Unix "shebang" line) may contain text in one of these formats:

```
-- mode: modename --
```

```
-- modename --
```

```
-- tab-width: number --
```

```
-- mode: modename; tab-width: number --
```

Other characters may appear before or after each possibility above; typically there would be commenting characters, so a full line might read `/* -- mode: shell -- */`. The first two examples set that buffer to the specified mode name, such as Perl or VBasic or C, by running a command or function named *modename*-mode if one exists. (A mode name of “C++” makes Epsilon use the C++ submode of C mode.) The third example sets the width of a tab character for that buffer.

In more detail, between the `--` sequences may be one or more definitions, separated by `;` characters. Spacing and capitalization are ignored throughout. Each definition may either be a mode name alone, or a setting name followed by a colon `:` and a value.

The setting names recognized are “mode”, as another way to specify the mode; “tab-size” or “tab-width” to set the buffer’s tab size, or “margin-right” or “fill-column” to set the buffer’s right margin. (The names `tab-size` and `margin-right` reflect the names of the Epsilon variables they set; the names “tab-width” and “fill-column” are more compatible with other programs, and recommended if users of other editors may edit the files.)

Similarly, you can use either “auto-fill-mode” or “fill-mode” to set whether Epsilon should break lines as you type, and either “indent-with-tabs” or “indent-tabs-mode” to set whether indenting should use tab characters in addition to spaces. The latter name, in each case, is the more compatible one. Also, you can write “nil” instead of 0 to turn off a setting, again for compatibility.

Epsilon also recognizes “compile-command” for use with the `compile-buffer` command; see page 154 for details. And it recognizes “coding” to indicate the file’s Unicode encoding (but the `detect-encodings` variable may override this).

It recognizes “write-line-translate” as a way to set the style of line translation for the file after it has been read in; this is useful as a directory-wide setting, to permit files to be auto-detected when read, but forced into a consistent format when written. The recognized value names for this setting are: “windows” (or equivalently “dos”), “binary”, “unix”, “mac”, and “auto”. See page 122 for details.

Another syntax for normal file variables only appears at the end of a file, starting within the last 3000 characters (but see the `file-variables-search-limit` variable). It looks like this:

```
Local Variables:
mode: modename
tab-size: number
End:
```

The first line shown is required, and the last is recommended for compatibility; inside are the settings, one per line. Each line may have additional text at the start and end of each line (so it will look like a comment in the file’s programming language). The “coding” file variable doesn’t use this alternative syntax; any specified encoding must be on the first line only.

When using this syntax, (or when using an `.epsilon_vars` file, as the next section explains), Epsilon also permits a file to set any numeric Epsilon variable defined as buffer-specific. These include `auto-fill-indents`, `auto-fill-mode`, `c-indent`, and many more. Such variables may also be set based on another such variable’s value:

```
c-indent: 6
c-contin-offset: c-indent
c-label-indent: c-indent * 3
c-case-offset: c-indent * -1
```

Values for numeric file variables use a simple syntax. They may be either a number (perhaps with a 0x, 0o, or 0b prefix to indicate its base is hex, octal or binary, respectively), or else the name of an Epsilon variable that could be used as a file variable, or else two of these with an operator between them (+, -, *, or / only). (Either a number or a variable name may have a - prefix to negate it.) Those file variable names that Epsilon recognizes only for compatibility (fill-column and tab-width, for instance) may not be used with these expressions.

Bits in the variable `use-file-variables` enable scanning for file variables of different sorts.

Directory-wide File Variables

You can put file variables in a special file named `.epsilon_vars`. Such settings apply to all files in its directory. Or you can create a file named `.epsilon_vars` in your Epsilon customization directory, and its settings will be applied first, before any directory-specific file.

If there's no `.epsilon_vars` file in a directory, you can tell Epsilon to search up the directory tree until it finds one in a parent directory. A bit in the `use-file-variables` variable enables this.

In an `.epsilon_vars` file, lines starting with # are comments. It contains one or more sections. Within each section, settings in it appear one per line, with a setting name, a colon, and a value. Each section begins with a line that says which files or modes it affects:

```
# Special settings for this directory.
Extensions: .r*
mode: Perl

Modes: Perl|Python
tab-size: 3

Modes: C
tab-size: 5
indent-tabs-mode: nil
```

The Modes, Extensions, Filenames, and Basenames lines use a file wildcard pattern. It can use | for alternation, ? to match a single character or * to match any number, or character ranges like [a-z]. Epsilon will apply the settings in the section that follows only if the original file's extension, mode, full pathname, or basename matches the pattern. This example says that all files with an extension like .r or .rxx or .ram in that directory should use Perl mode, and sets the tab size to 3 for Perl or Python files, and 5 for C files, also turning off using Tab characters to indent.

Here's another example you might find in a system-wide `.epsilon_vars` file.

```
# Make every .inc file in the hierarchy below be in PHP mode, instead
# of auto-detecting.
Filenames: /home/bob/hats/*.inc
mode: PHP

# Any file matching one of these names, in any directory, is really
# a makefile.
Basenames: buildfile*|build[1-2]*|*build.dat
mode: makefile
```

Epsilon decides which sections to use before applying the settings, so an .rx file forced to Perl mode by the above example file won't get a tab size of 3 unless you add a `tab-size: 3` line to its Extensions section. Also note that "mode:" sets a file's mode; "Modes:" begins a section for a specific mode. File variables in an individual file take precedence over those in an .epsilon_vars file, and those in a directory-specific .epsilon_vars file take precedence over those in a system-wide file.

Vi/Vim File Variables

Epsilon also supports a few file variables using an alternative syntax used by the Vi/Vim family of editors.

Each such setting line (which Vi/Vim documentation refers to as a "modeline") must appear within five lines of the start or end of the file. They begin with a space, then the word "vi" (or alternatively, "vim" or "ex"), followed by a colon. One format follows this with the word "set" or "se", then a series of settings separated by spaces and terminated by a colon; other text on the line can surround this. The other format omits "set" and uses a series of settings separated by spaces or colons and terminated by the end of the line.

Here are examples of each type:

```
/* vim: set textwidth=65 tabstop=8 sts=3 noexpandtab: */
; vi: tw=70 ts=4:softtabstop=2 et
```

The Vi/Vim settings Epsilon recognizes are:

Vi Setting Name	Vi Synonym	Epsilon Equivalent
<code>textwidth=val</code>	<code>tw=val</code>	<code>margin-right=val</code>
<code>tabstop=val</code>	<code>ts=val</code>	<code>tab-size=val</code>
<code>shiftwidth=val</code>	<code>sw=val</code>	<code>tab-size=val</code>
<code>softtabstop=val</code>	<code>sts=val</code>	<code>soft-tab-size=val</code>
<code>expandtab</code>	<code>et</code>	<code>indent-with-tabs=0</code>
<code>noexpandtab</code>	<code>noet</code>	<code>indent-with-tabs=1</code>

4.8.4 Internet Support

Epsilon for Windows or Unix has several commands and facilities that make it easy for you to edit files on other computers using the Internet.

The `find-file` and `dired` commands, as well as a few others, understand Internet URLs. If you provide the URL `ftp://user@example.com/myfile.c` to a file-reading command like `find-file`, Epsilon will engage in an FTP interaction to download the file and display it in a buffer. All of the Internet activity happens in the background, so you don't have to wait for the file to download before continuing with your work. In fact, the file appears in the buffer as it downloads (syntax highlighted if appropriate), so you can be editing the beginning of a large file while the rest of it downloads.

Saving a file in such a buffer, or writing a buffer to a file name that starts with `ftp://`, will cause Epsilon to send the file to the remote computer. Upload and download status is indicated in the mode line, and there's also a `show-connections` command (on `Ctrl-Alt-c`) that shows the status of all Internet activities and buffers. As in `bufed`, you can select a buffer and press `(Enter)` to switch to it, or press `(Escape)` to remain in the current buffer. Use the `kill-process` command to cancel an FTP transfer or Telnet session (see below) in progress in the current buffer.

FTP and SCP URLs (the latter described in the next section) work with `dired` also, so if you do a `dired` (or a `find-file`) on `ftp://user@example.com`, you'll get a directory listing of the files on the remote machine `example.com`, in a familiar `dired` context. `Dired` knows how to delete and rename remote files, and sort by size, date, file name or extension. To make `Epsilon` work with certain host computers (systems running VMS, for example), you may need to set the variables `ftp-ascii-transfers` or `ftp-compatible-dirs`; see the descriptions of those variables in the online help. Other systems may require you to set the variable `ftp-passive-transfers`.

The `telnet` command lets you connect to a command shell on a remote computer. The `ssh` command described in the next section provides the secure equivalent. Each creates a buffer that works much like the `Epsilon` process buffer, except the commands you type are executed on the remote machine. Provide a numeric prefix argument and `telnet` will connect on the specified port instead of the default port. Or use the syntax `hostname:port` for the host name to specify a different port. You can either use the `telnet` command directly, or specify a `telnet: URL to find-file`. (`Epsilon` ignores any username or password included in the URL.) Typing `Ctrl-c Ctrl-c` in `telnet` or `ssh` buffers sends an interrupt signal to the remote system, aborting the current program.

In a `telnet` buffer, the `telnet-interpret-output` variable controls whether `Epsilon` interprets certain ANSI color-setting escape sequences and similar things. (The `ssh-interpret-output` variable is the equivalent for `ssh`.) `Epsilon` also looks for password requests from the remote system, using the `recognize-password-pattern` variable, so it can hide the password as you type it. Also see the `recognize-password-prompt` variable, and the `send-invisible` command.

Normally `Epsilon` doesn't send a line in a `telnet` or `ssh` buffer until you press `(Enter)`. Type `Ctrl-u (Enter)` to send a partial line immediately.

As in a concurrent process buffer, you can press `Alt-p` or `Alt-n` to access a `telnet` or `ssh` buffer's command history. With a numeric prefix argument, these keys show a menu of all previous commands. You can select one to repeat.

If you specify an `http: URL to find-file` (for example, `http://www.lugaru.com`), `Epsilon` will use the HTTP protocol to retrieve the HTML code from the given location. The HTML code will appear in an appropriately named buffer, syntax highlighted. Header information for the URL will be appended to a buffer named "HTTP Headers".

You can tell `Epsilon` to send its requests by way of a proxy by setting the variables `http-proxy-server`, `http-proxy-port`, and `http-proxy-exceptions`. You can tell `Epsilon` to identify itself to the server as a different program by setting `http-user-agent`, or set `http-force-headers` to entirely replace `Epsilon`'s HTTP request with another, or to add other headers. The `http-log-request` variable makes `Epsilon` copy the entire request it sends to the HTTP Headers buffer.

If you specify an `https: URL to find-file`, `Epsilon` will try to run the `curl` or `wget` programs, if they are installed, to retrieve the URL's contents and insert it in an appropriately-named buffer. Windows users can install either of these programs via the Cygwin system. Mac OS systems typically have `curl` installed by default, but users can install either via the Homebrew system if necessary. Linux and FreeBSD users can install via their system's package installers.

The `Alt-e` and `Alt-(Down)` keys in `find-file` come in handy when you want to follow links in an HTML buffer; see page 30 for information on `Alt-e` and page 28 for information on `Alt-(Down)`. Also see the `find-linked-file` command on `Ctrl-x Ctrl-l`.

The command `view-web-site` on `Shift-F8` searches for the next URL in the buffer. It prompts with that URL, and after you modify it if necessary, it then launches an external browser on the URL. The `view-lugaru-web-site` command launches a browser and points it to Lugaru's web site. `Epsilon` for Unix uses a shell script named `goto_url` to run a browser. See page 41. `Epsilon` for Windows uses the system's default browser.

If you run a compiler via telnet or a similar process in an Epsilon buffer, you can set up the next-error command on Ctrl-x Ctrl-n so that when it parses a file name in an error message, it translates it into a URL-style file name that Epsilon can use to access the file. To do this, you'll need to write your own `telnet_error_converter()` subroutine in EEL. See the sample one in the Epsilon source file `epsnet.e` for details.

Summary:	Ctrl-Alt-c	show-connections
	Telnet mode only: Alt-n	process-next-cmd
	Telnet mode only: Alt-p	process-previous-cmd
		telnet
		telnet-mode
		view-web-site
		view-lugaru-web-site

Secure Shell and SCP Support

Besides recognizing `ftp://` URLs as described in the previous section, Epsilon also recognizes `scp://` URLs, which may be used for secure file transfers. With scp support, you can read or write files using an `scp://` URL, navigate the remote system's directory tree using `direcd`, mark files for copying between the local and remote systems, use `grep` or `file-query-replace` to search and replace on multiple remote files, and use file name completion.

Epsilon also recognizes `ssh://` URLs to connect securely to a command shell on a remote computer, providing a secure alternative to the `telnet` command. Epsilon's `ssh` command works similarly to the `ssh://` URL. Use the syntax `username@hostname` to connect as a user other than the default one. The `ssh-interpret-output` variable controls how Epsilon interprets ANSI escape sequences and similar in an `ssh` buffer.

The `scp` and `ssh` features work by running certain external programs which must be installed and working. Epsilon's `ssh` command depends on an external `ssh` program, while its `scp` features run a program named `sftp`. On macOS these are normally preinstalled. For Linux or FreeBSD, you may need to install the appropriate `ssh` package for your distribution. For Windows, the Cygwin system contains appropriate clients. Run the Cygwin installer from the Cygwin website <https://www.cygwin.com> and install Cygwin's `openssh` package from the `net` section. Also ensure Cygwin's `bin` directory is on your `PATH`, and (on Windows 10) that it appears before the Windows system directory (which contains programs of the same name as the Cygwin ones but which don't work well with Epsilon). See below for a step-by-step setup guide. (On Windows, it's also possible to use alternative clients like PuTTY instead of Cygwin programs. See "Windows-specific Configuration" below for more on PuTTY.)

With `scp/ssh` support, Epsilon doesn't remember your password or passphrase. Unless you've set up the `ssh` and `scp` programs so they don't require a password (see below for a step-by-step guide), Epsilon will ask for it each time it must start a new `sftp` helper program (for instance, when you begin a second file operation before the first has completed). If you prefer to type your secure passphrase once and have multiple connections use it, you can set up an `ssh-agent` program, along with public key authentication. The agent will remember your credentials and provide them as required to any `sftp` or `ssh` instance. You can even set your credentials to expire after a certain period of time if you wish. Refer to the manual page for the `ssh-agent` program to set this up. Windows users should also see the section below on Windows-specific configuration.

CUSTOMIZATION OPTIONS FOR ALTERNATIVE CLIENTS

If you're not using the usual external `ssh` and `sftp` programs, you'll need to set various variables to tell Epsilon how to run your alternative programs.

The variable `ssh-template` tells Epsilon how to build a command line for invoking the external `ssh` program when a specific user name appears before the host name. If no user name was specified, it uses `ssh-no-user-template`. See the descriptions of these variables for their format. Both these variables normally include the `%c` sequence, which under Windows makes Epsilon run the command named by the `ssh-command-windows` variable. There are also numerous variables whose names start with `sftp-` that may be used to configure Epsilon to work with alternatives to the `sftp` program.

Some very old `sftp` programs use a different command syntax for listing files; if you have trouble, try setting the `scp-client-style` variable to 2 to make Epsilon use old-style `sftp` commands. You may have to modify `scp-list-flags` too.

WINDOWS-SPECIFIC CONFIGURATION OPTIONS

As explained above, using an `ssh-agent` program along with public key authentication lets you type your secure passphrase once and have multiple connections use it. The agent must provide some settings that are passed on to the `sftp` or `ssh` clients it runs via environment variables. For Windows users running Cygwin, one option is to start Cygwin's bash shell, run the command `eval `ssh-agent``, run the `ssh-add` command, and then run Epsilon from that same shell. Or you can use the `run-ssh-agent.bat` file included in Epsilon's `bin` subdirectory to run an `ssh` agent. The comments in that file explain how to run `ssh-agent` through it, so it creates a `load-ssh-agent` batch file that loads agent settings into the environment, and how to set Epsilon variables so Epsilon invokes `load-ssh-agent` when starting `ssh` or `scp` sessions.

To make Epsilon work with the Windows `ssh` client PuTTY instead of the recommended Cygwin clients, use these settings:

```
scp-windows-sftp-command  psftp
ssh-command-windows       plink
scp-client-style           2
```

Be sure to install PuTTY's `psftp` and `plink` programs along with the base PuTTY installation. With PuTTY, certain features like file name completion won't be available.

TYPICAL CONFIGURATION WALKTHROUGH

Here are some typical steps required to get Epsilon running with secure shell and `scp` support. We'll assume the goal is to connect from a computer named `mypc` to a computer named `otherbox` which runs some kind of Unix-derived operating system, and where you log in using the name "myname".

Part 1: Set up the basic connection

- If `mypc` runs Windows, run the Cygwin installer from <https://www.cygwin.com> and have it install its `ssh` package, which includes the `ssh` and `sftp` programs. (Note that Windows 10 includes `ssh` and `sftp` clients with the same names as the Cygwin clients, but which don't work well with Epsilon. Either ensure the Cygwin clients appear first on your `PATH`, or modify the variables `ssh-command-windows` and `scp-windows-sftp-command` to include the full path to the Cygwin clients.) On non-Windows systems, install `ssh` if it's not installed already.
- Set up `otherbox` to run an `ssh` server, which may require installing and/or enabling the server. Test that `mypc` can access `otherbox` by name (for example, by typing "`ping otherbox`" at a command prompt).
- You should now be able to run `ssh` from a command prompt on `mypc` to connect to `otherbox`, supplying your `otherbox` user name on the command line if it differs from your `mypc` user name: "`ssh myname@otherbox`", and typing your password on `otherbox`. At this point, you can edit over `ssh` by typing `scp://myname@otherbox` at a file name prompt in Epsilon (or `scp://otherbox` if your names match). But you'll need to type your password for each editing session.

Part 2: Set up an automatic connection

- If your security needs permit, you may want to set up an automatic connection, where anyone signed into your account on mypc would also be able to access otherbox, with no further authorization required. To do this, first run `ssh-keygen` on mypc to generate a pair of files in your `~/.ssh` directory containing a key pair, and enter an empty passphrase. (We'll assume you used the default file name `"id_rsa"`.)
- On mypc, run `"scp ~/.ssh/id_rsa.pub otherbox:~/.ssh/authorized_keys"` to permit connections from mypc to otherbox using the key pair you just created. (If you already have an `authorized_keys` file on otherbox, perhaps because you connect to it over ssh from other systems, add the line from `id_rsa.pub` to its end, don't overwrite it.) Test this by running `"ssh myname@otherbox"`; it should connect without any prompt for a password.
- You should now be able to edit over ssh by typing `scp://myname@otherbox` at a file name prompt in Epsilon, and immediately see a file listing.

Optional: If you'd like to use a shorthand name for otherbox, or skip typing `"myname@"` each time, or if you chose a non-default key pair name when you ran `ssh-keygen`, you'll need an ssh configuration file. On mypc, edit `\cygwin64\home\YourWindowsUserName\.ssh\config` (or `~/.ssh/config` on non-Windows systems) and add a section like this:

```
Host oth
  Hostname otherbox
  User myname
  IdentityFile ~/.ssh/mynewkey
```

This defines `"oth"` as a short name for `"otherbox"`. (You can skip this and just use a `Host` line if you don't need a shorter name.) It says your login name on otherbox is `"myname"`, so you don't have to include it when connecting. You can omit this line if your login name is the same on both systems. And it specifies a different private key file name. You can include other ssh configuration options here; for example, macOS users may want to include `"UseKeychain yes"` to have ssh get passphrases from your keychain.

PER-SYSTEM SETTINGS

It's possible to set up Epsilon to use one set of variables for one remote system and a different one for others. To enable this, before checking for a variable such as `scp-run-helper-template`, Epsilon constructs a new variable name by adding the host name of the remote system to its end. For instance, if you try to access `www.example.com`, Epsilon first looks for a variable named `scp-run-helper-template-www-example-com`; if there's a variable by that name, Epsilon uses it instead of the usual one. (Epsilon constructs the variable name from a host name by replacing each non-alphanumeric character with a `-`.) It does this for each of its `scp` and `ssh` variables.

FILE PROMPTS IN SSH AND TELNET BUFFERS

When you're editing a local file, commands like `find-file` prompt, by default, with that file's directory name, so you can easily edit another file in that same directory. Begin typing an absolute pathname to another directory, and the old directory name is automatically deleted. In the same way, when you use commands like `find-file` from an ssh or telnet buffer, Epsilon prompts with a corresponding `scp: URL` so you can access files on that same system, and in your current directory on that remote system.

This feature requires Epsilon to parse the prompts coming from the remote system, in order to know what directory you're using. Out of the box, it's designed to work with some typical prompts you might get when connecting to a remote Unix-like system. You can configure it, though, on a per-host basis.

Epsilon uses the `net-prompt-pattern` variable to match a prompt that contains a directory name. The default setting works with a typical Unix shell prompt setting (PS1) like `[\u@\h \w]$` or `[\u@\h \w] \!`. A prompt-matching pattern should be written so that its first parenthesized grouping matches the part of your prompt containing your current directory on the remote system.

Alternatively, you can change your prompt to include a special sequence of characters at the end containing your current directory, which Epsilon will notice and remove. This is how Epsilon's concurrent process retrieves current directory information on Unix-based systems, and it's more reliable because it avoids shell settings that might abbreviate the full working directory when it's long. Epsilon includes scripts named `epsilon.sh` and `epsilon.csh` (for Bourne-style and Csh-style shells) to accomplish this, so you could configure your remote system's login script to run them only for remote connections (by checking the `TERM` environment variable, for example).

If your connection to a Unix-based host sets the `TERM` variable to something other than "dumb", you may see odd line breaks in echoed text. A command like "stty columns 65000" should prevent that, and again, this can be done conditionally in the remote system's login script.

You can set up a separate `net-prompt-pattern` variable for each host you use, by defining a variable with a name like `net-prompt-pattern-example-com`, as explained in the previous section. You can also write a custom per-host EEL function to translate the retrieved directory into a URL, by defining a function with a name like `net_prompt_get_dir_helper_example_com()`. See the definition of `net_prompt_get_dir_helper()` in `epsnet.e`.

You can configure various aspects of remote directory retrieval by setting bits in the `net-prompt-get-directory` variable. Like other `net` variables, you can create host-specific versions of the variable. Omit the 1 bit in the original `net-prompt-get-directory` variable to turn off remote directory retrieval for all hosts; omit it in a per-host variable to disable it for just that host.

USING ANCIENT HOSTS

If you must use a very old version of `ssh` that lacks an `sftp` program, or connect to a system that doesn't support `sftp`, or you want to use an `ssh` replacement that lacks `sftp`, it's possible to set up Epsilon to run its own helper program on the remote system.

To do this, copy the C language source code file `epsilon-xfer-helper.c` included in Epsilon's source directory to the remote system, compile it with "make `epsilon-xfer-helper`" or similar, and install in an accessible location. It may be compiled on most Unix systems, or, for Windows, using the Cygwin environment. Next, check that you can run the helper program remotely, with a command line like

```
ssh -l username hostname epsilon-xfer-helper
```

It should print a greeting line and await a command. Type `^C` or press `<Enter>` to make it exit. You may need to edit the Epsilon variable `scp-run-helper-template` to include the path to the helper program, or if you use a different `ssh` program. For instance, if you use an `ssh` client "oldssh" that lacks an `sftp` program, set it to "oldssh `%u@%h /path/to/epsilon-xfer-helper`" or similar. (Epsilon uses the above variable when the `scp://` url includes a user name, and the `scp-run-helper-no-user-template` variable when it does not.)

To tell Epsilon to use `epsilon-xfer-helper` commands, not `sftp` commands, set the `scp-client-style` variable to 1. Using the helper program enables a few minor features that the `sftp` program doesn't currently support, like using `~` to indicate home directories, or copying a remote file to a different location on the remote system (`sftp` can rename remote files but not copy them).

When you don't use `sftp`, Epsilon must run a separate program for each file transfer. By default it uses the `scp` program. The variable `scp-read-file-template` tells Epsilon how to transfer a file from the remote system to a local file, and `scp-write-file-template` does the opposite. There are separate versions of these variables for when no user name is included, named

`scp-read-file-no-user-template` and `scp-write-file-no-user-template`. Change these variables to use a different program for copying files when you don't use sftp.

Summary:

	ssh
	ssh-mode
Ssh mode only: Alt-n	process-next-cmd
Ssh mode only: Alt-p	process-previous-cmd

URL Syntax

In Epsilon, URLs must start with `ftp://`, `http://`, `https://`, `scp://`, `ssh://`, or `telnet://`. (If you omit the service name, the `ftp:` part, Epsilon for Windows will pass the file name to Windows as a UNC-style network file name.)

For some services, you can specify a user name, password, or port number using the URL syntax of `service://username:password@hostname:portnumber/filepath`. (Ftp and http recognize all three, telnet recognizes only a port number, and scp recognizes only a user name.)

If you include a user name in an ftp or http URL but omit the `:password` part, Epsilon will prompt for one (and will make sure the password does not appear in your state file, session file, or similar places). But if you include a password in your URL, note that it may be saved in Epsilon's session file or similar places.

If you omit the `username:password@` or `username@` part entirely in an ftp URL, Epsilon uses the user name "anonymous" and the password specified by the `anon-ftp-password` variable (default: `EpsilonUser@unknown.host`). You can set this to your email address if you prefer.

You can also use Emacs-style syntax for specifying remote file names: `/username@hostname:filepath`. Epsilon will behave as if you had typed the corresponding URL.

In `ftp://` URLs, Epsilon treats a file name following the `/` as a relative pathname. That is, `ftp://user@example.com/myfile` refers to a file named `myfile` in the user's home directory. Put two slashes, as in `ftp://user@example.com//myfile`, to refer to `/myfile` in the root directory. You can type `\` instead of `/` in any URL and Epsilon will substitute `/`.

If you type the name of a local directory to the `find-file` command, `find-file` will run the `dired` command on it. With `ftp://` URLs, `find-file` won't always know that what you typed is a remote directory name (as opposed to a file name) and might try to retrieve the URL as a file, leading to an error message like "Not a plain file". End your URL with a `/` to indicate a directory name.

4.8.5 Unicode Features

This section explains how to use Epsilon to edit text containing non-English characters such as `ê` or `å`.

Epsilon supports Unicode, as well as many 8-bit national character sets such as ISO 8859-1 (Latin 1).

In Unix, Unicode support is only available when Epsilon runs under X11, and when a font using the `iso10646` (Unicode) character set is in use. See <https://www.lugaru.com/links.html#unicode> for Unicode font sources. Epsilon includes a shell script named `get_unicode_core_x11_fonts` that can install Unicode-based fonts in various sizes. By default it's in `/opt/epsilon14.00/bin` on Linux and FreeBSD, and `/Applications/Epsilon 14.00.app/Contents/Resources` on macOS. Under Unix, Epsilon displays all characters using a glyph width determined by the widest character in the font.

Epsilon for Windows shows characters in a font using their specified individual widths, but only in a full-width window. If you instead create side-by-side windows, Epsilon will ignore the special width rules

of zero-width characters and extra-wide characters, among other things, putting each character into a same-width cell. Text with such characters should be edited in full-width windows for the best display. (See the `change-show-spaces` command to make zero-width characters visible.)

To enable Unicode display in Epsilon for Windows Console, see the `console-ansi-font` variable. Also see page 124 for more information on the DOS/OEM encoding used by default in the Windows Console version.

In this release, Epsilon doesn't display Unicode characters outside the basic multilingual plane (BMP), or include any of the special processing needed to handle complex scripts, such as scripts written right-to-left. Each character outside the BMP is handled as a pair of surrogate characters. While Epsilon cannot display their glyphs, the `show-point` command will report the Unicode character name of the one at point.

Epsilon knows how to translate between its native Unicode format and dozens of encodings and character sets (such as UTF-8, ISO-8859-4, or KOI-8).

Epsilon autodetects the encoding for files that start with a Unicode signature ("byte order mark"), and for many files that use the UTF-8 encoding. To force translation from a particular encoding, provide a numeric argument to a file reading command like `find-file`. Epsilon will then prompt for the name of the encoding to use. Press "?" when prompted for an encoding to see a list of available encodings. The special encoding "raw" reads and writes 8-bit data without any character set translation.

Epsilon uses the buffer's current encoding when writing or rereading a file. Use the `set-encoding` command to set the buffer's encoding.

The `unicode-convert-from-encoding` command makes Epsilon translate an 8-bit buffer in a certain encoding to its 16-bit Unicode version. The `unicode-convert-to-encoding` command does the reverse.

You can add a large set of additional converters to Epsilon by downloading a file. Mostly these converters add support for various Far East languages and for EBCDIC conversions. See <https://www.lugaru.com/encodings.html> for details.

Internally, buffers with no character codes outside the range 0–255 are stored with 8 bits per character; other buffers are stored with 16 bits per character. Epsilon automatically converts formats as needed.

The `detect-encodings` variable controls whether Epsilon tries to autodetect certain UTF-8 and UTF-16 files. The `default-read-encoding` variable says which encoding to use when autodetecting doesn't select an encoding. The `default-write-encoding` variable sets which encoding Epsilon uses to save a file with 16-bit characters and no specified encoding, in a context where prompting wouldn't be appropriate such as when auto-saving.

See the `insert-ascii` command on page 57 to type arbitrary Unicode characters, and the `show-point` command to see what specific characters are present (if the current font doesn't make that clear enough).

4.8.6 Printing

The `print-buffer` command on Alt-F9 prints the current buffer. If a region is highlighted on the screen, the command prints just that region. The `print-region` command on Shift-F9 always prints just the current region, whether or not it's highlighted.

Under Windows, the printing commands display the familiar Windows print dialog. From this dialog, you can select a different printer, select particular pages to print, and so forth. The `print-setup` command lets you select a different printer without printing anything, or set the margins. Invoke the printing commands with a numeric prefix argument to skip the print dialog and just print with default settings. The `print-buffer-no-prompt` command also skips the print dialog and uses default settings.

You can change the font Epsilon for Windows uses for printing with the `set-printer-font` command. See page 112 for more information.

By default, Epsilon for Windows will print in color on color printers, and in black & white on non-color printers. You can set the `print-in-color` variable to 0, if you don't want Epsilon to ever print in color, or to 2 if you want Epsilon to attempt to use colors even if the printer doesn't appear to be a color printer. (Some printers will substitute shades of grey.) The default value, 1, produces color printing only on color printers.

If you have a color printer, and want to use a different color scheme when printing than you do for screen display, set the variable `print-color-scheme` to the name of the color scheme Epsilon should use for printing.

Epsilon for Windows prints a heading at the top of each page. You can set the `print-heading` variable (which see) to control what it includes. By default it prints the file name, page number, and current date.

You can set the variable `print-line-numbers` nonzero if you want Epsilon to include line numbers, or set `print-doublespaced` if you want Epsilon for Windows to skip alternate lines. (To display line numbers on the screen, not when printing, see the `draw-line-numbers` variable.)

In non-Windows environments, the printing commands prompt for the device name of a printer. They then write the text to that device name. If you want Epsilon to run a program that will print the file, you can do that too. See the description of the `print-destination` variable in the online help. (For Unix, see `print-destination-unix`, which by default runs the `lpr` program to print a file.) If you want Epsilon for Windows to run a program in order to print a file, bypassing the Windows print dialog, you can set `want-gui-printing` to zero.

By default, Epsilon converts tabs to spaces in a copy of the buffer before printing it. Set the variable `print-tabs` to one if you want Epsilon to print the file just as it is, including the tab characters.

Summary:	Alt-F9	<code>print-buffer</code>
	Shift-F9	<code>print-region</code>
		<code>print-setup</code>

4.8.7 Extended file patterns

This section describes Epsilon's extensions to the rules for wildcard characters in file names. You can specify more complicated file name patterns in Epsilon than Windows or Unix normally allow, using the wildcard characters of square brackets `[]`, commas, semicolons, and curly braces `{}`. Epsilon also lets you use the `*` and `?` characters in more places. These patterns work in the `grep` command, the `dired` command, and in all other places where file name wildcards make sense. (They don't work with Internet URLs, though.)

First, you can put text after the standard wildcard character `*` and Epsilon will match it. In traditional DOS-style patterns, the system ignores any text in a pattern between a `*` and the end of the pattern (or the dot before an extension). But in Epsilon, `ab*ut` matches all files that start with `ab` and end with `ut`. The `*` matches the dot character in file names, so the above pattern matches file names like `about` as well as `absolute.out`. (Use `ab*ut.` to match only files like the former, or `ab*.*ut` to match ones like the latter.) Epsilon for Windows also follows the DOS/Windows convention that a file pattern ending with a `.` matches files with no extension.

Instead of `?` to match any single character (except dot, slash, or backslash), you can provide a list of characters in square brackets (similar to the regular expression patterns of searching). For example, `file[0123456789stuvw]` matches `file4`, `file7`, and `files`, but not `filer`. Inside the square brackets, two characters separated by a dash represent a range, so you could write the above pattern as `file[0-9s-w]`. A caret character `^` just after the `[` permits any character but the listed ones, so

`fil[tm]er` matches all the files that `fil?er` matches, except `filter` and `filmer`. (To include a dash or `]` in the pattern, put it right after the `[` or `^`. The pattern `[^-]` matches all characters but `-` and `]`.)

You can use `?` and `*` (and the square bracket syntax) in directory names. For example, `\v**.bat` might match all `.bat` files in `\virtmem` and in `\vision`. Because a star character never matches backslash characters, it would not match `\vision\subdir\test.bat`.

The special directory name `**` matches any number of directory names (including zero). You can use it to search entire directory trees. For example, `***.txt` matches all `.txt` files on the current drive. The pattern `**\include*.h` matches all `.h` files inside an `include` directory, looking in the current directory, its subdirectories, and all directories within those. A pattern ending in `**` matches all files in that hierarchy. You can set the `file-pattern-ignore-directories` variable to have Epsilon skip over certain directories when expanding `**`.

The simplest new file pattern character is the comma. You can run `grep` on the file pattern `foo,bar,baz` and Epsilon will search in each of the three files. You can use a semicolon in place of a comma, if you want.

A segment of a file pattern enclosed in curly braces may contain a sequence of comma-separated parts. Epsilon will substitute each of the parts for the whole curly-brace sequence. For example, `\cc\include\c*t.{bat,txt}` matches the same files as `\cc\include\c*t.bat,\cc\include\c*t.txt`. A curly-brace sequence may not contain another curly-brace sequence, but may contain other wildcard characters. For example, the pattern `{,c*}*. {txt,bat}` matches `.txt` and `.bat` files in the current directory, or in any subdirectory starting with `"c"`. The brace syntax is simply a shorthand for the comma-separated list described above, so that an equivalent way to write the previous example is `*.txt,c**.txt,*.bat,c**.bat`. Epsilon breaks a complete pattern into comma-separated sections, then replaces each section containing curly braces with all the possible patterns constructed from it. You can use semicolons between the parts in braces instead of commas if you prefer.

To match file names containing one of the new wildcard characters, enclose the character in square brackets. For example, the pattern `abc[]]` matches the file name `abc}`. (Note that legal DOS file names may not contain any of the characters `[]`, `;`, but they may contain curly braces `}`. Other file systems, including Windows VFAT, Windows NT's NTFS, most Unix file systems, and OS/2's HPFS, allow file names that contain any of these characters.)

Use curly braces to search on multiple drives. `{c,d,e}:***.txt` matches all `.txt` files on drives `C:`, `D:`, or `E:`. Epsilon does not recognize the `*`, `?`, or `[]` characters in the drive name.

When a file name contains a literal brace character, a comma, or one of the other characters used for extended wildcard patterns, you can surround it in quotes (`"`) to tell Epsilon to treat it literally, not as a wildcard pattern. Or you can set the `file-pattern-wildcards` variable to disable the wildcarding function of specific characters. If your file names often contain commas, for instance, you may want to disable comma's wildcard function.

It's possible to make Epsilon ignore certain types of symbolic links (and similar Windows NTFS file system entities) when interpreting file patterns. For instance, you can keep a `**` pattern from matching symbolic links to directories, only matching actual directories. See the `file-pattern-rules` variable.

Under Windows, file pattern matching also matches on the names of NTFS streams, and on the server and share names of UNC files. You can restrict server name matching to particular domains to speed it up on large networks; see the `file-pattern-unc-domains` variable. You can also use the pattern `*`: to match Windows drives (see `drive-match-types` to set which types).

4.8.8 Directory Editing

Epsilon has a special mode used for examining and changing the contents of a directory conveniently. The `dired` command, bound to `Ctrl-x d`, asks for the name of a directory and puts a listing of the directory, similar

to what the DOS “dir” command produces (or, for Unix, “ls -lF”), in a special dired buffer. By default, dired uses the current directory. You can supply a file pattern, such as “*.c”, and only matching files will appear. The dired command puts the information in a buffer whose name matches the directory and file pattern, then displays the buffer in the current window. You can have multiple dired buffers, each displaying the result of a different file pattern.

You can also invoke dired from the find-file command. If you press `<Enter>` without typing any file name when find-file asks for a file, it does a dired on the current directory. If you give find-file a file name with wild card characters, it runs the dired command giving it that pattern. If you give find-file a directory name, it does a dired of that directory. (When using ftp:// URLs that refer to a directory, end them with `/`. See page 136 for details.)

You can use extended file patterns to list files from multiple directories. (See page 138.) If you use a file pattern that matches files in more than one directory, Epsilon will divide the resulting dired buffer into sections. Each section will list the files from a single directory. Epsilon sorts each section separately.

While in a dired buffer, alphabetic keys run special dired commands. See the next section on page 140 for a complete list.

The quick-dired-command command on `Alt-o` is like running a dired on the current file, then executing a single dired command and discarding the dired buffer. It provides a convenient way of performing various simple file operations without running dired. It prompts for another key, one of C, D, M, G, !, T, V, or @. Then it (respectively) copies, deletes, or renames the current file, changes Epsilon’s current directory to the one containing that file, runs a command on the file, shows the file’s properties dialog, views it using associations, or toggles whether it’s read-only. `Alt-o +` creates a new directory, prompting for its name. `Alt-o .` displays a dired of the current file. `Alt-o a` lets you set the file’s attributes or permission bits. `Alt-o f` views its folder in File Explorer on Windows, the Finder on macOS, or corresponding programs on Unix. The other keys are similar to their corresponding dired subcommands; see the next section for more details. (The T option is only available in Epsilon for Windows.)

By default, Epsilon records dired buffers in its session file and recreates them the next time you start Epsilon, except for remote direds that use a URL. See the variables `session-restore-directory-buffers` and `session-restore-max-directories`.

The locate-file command prompts for a file name and then searches for that file, using dired to display the matches. In Windows, it searches for the file on all local hard drives, skipping over removable drives, CD-ROM drives, and network drives. On Unix, it searches through particular parts of the directory hierarchy specified by the `locate-path-unix` variable.

The list-files command also takes a file pattern and displays a list of files. Unlike dired, its file list uses absolute pathnames, and it omits the file’s size, date, and other information. It provides just the file names, one to a line. The command also doesn’t list directory names, as dired does. The command is often useful when preparing response files for other programs.

Summary:	<code>Ctrl-x d</code>	dired
	<code>Alt-o</code>	quick-dired-command
		list-files

Dired Subcommands

This section lists the subcommands you can use when editing a dired buffer (see page 139). You run most dired commands by pressing plain letters. All other keys still invoke the usual Epsilon commands.

The N and P commands go to the next and previous files, respectively.

The E, <Space>, and <Enter> keys let you examine the contents of a file. They invoke the find-file command on the file, making the current window display this file instead of the dired buffer. To conveniently return to the dired buffer, use the select-buffer command (Ctrl-x b). Press <Enter> when prompted for the buffer name and the previous buffer shown in the current window (in this case, the dired buffer) will reappear.

When applied to a subdirectory, these keys invoke another dired on that directory, using the name of the directory for that dired buffer. If you have marked files for deletion or copying, and you run a dired on the same directory, the markings go away.

The ‘.’ or “~” keys invoke a dired on the parent directory of the directory associated with the current dired buffer.

To set Epsilon’s current directory to the directory being displayed, press G (for Go). If the current line names a directory, Epsilon will make that be the current directory. If the current line names a file, Epsilon will set the current directory to the one containing that file.

Press D to flag a file that you wish to delete. Epsilon will mark the file for deletion by placing a ‘D’ before its name. (You may delete empty directories in the same way.) Press C or M to select files for copying or moving (renaming), respectively. Epsilon will mark the files by placing C or M before their names. The U command unmarks the file on the current line, removing any marks before its name.

The X command actually deletes, copies, or moves the marked files. Epsilon will list all the files marked for deletion and ask you to confirm that you want them deleted. If any files are marked for copying or moving, Epsilon will ask for the destination directory into which the files are to be copied or moved. If there is only one file to copy or move, you can also specify a file name destination, so you can use the command for renaming files. (In this case, Alt-g will copy the original file name so you can edit it.) Epsilon prompts for a single destination for all files to be copied, and another for all files to be moved.

If you try to delete a read-only file, Epsilon will prompt first; see the `dired-confirmation` variable to change this. If you try to delete a non-empty directory, Epsilon will similarly ask for confirmation before deleting the entire directory hierarchy. Similar prompts occur if you try to overwrite an existing local file when copying or moving a file.

There are a few specialized commands for renaming files. Press Shift-L to mark a file for lowercasing its name, or Shift-U for uppercasing. When you execute with X, each marked file will be renamed by changing each uppercase character in its name to lowercase (or vice versa). (Note that Epsilon for Windows displays all-uppercase file names in lowercase by default, so Shift-U’s effect may not be visible within Epsilon. See `preserve-filename-case`.)

Shift-R marks a file for a regular-expression replacement on its name. When you press X to execute operations on marked files, Epsilon will ask for a pattern and replacement text. Then, for each file marked with Shift-R, Epsilon will take the file name and perform the indicated regular expression replacement on it, generating a new name. Then Epsilon will rename the file to the new name. For instance, to rename a group of files like `dir\file1.cxx`, `dir\file2.cxx`, etc. to `dir2\file1.cpp`, `dir2\file2.cpp`, use Shift-R and specify `dir\(.*)\.cxx` as the search text and `dir2\#1\.cpp` as the replacement text. To rename some `.htm` files to `.html`, specify `.*` as the search text and `#01` as the replacement text.

By default, most files or directories that start with a period character `.` will be hidden. Pressing `-` toggles whether such files are hidden. The `dired-show-dotfiles` variable sets which files or directories are always shown regardless of this toggle. By default, dired entries for the current directory (`.`) and its parent (`..`) are always shown.

The `!` dired subcommand prompts for a command line, then runs the specified program, adding the name of the current line’s file after it. If the command line you type contains an `*`, Epsilon substitutes the current file name at that position instead of at the end. If the command line ends in a `&` character, Epsilon runs the program asynchronously; otherwise it waits for the program to finish.

The `+` command creates a new subdirectory. It asks for the name of the subdirectory to create.

The `R` command refreshes the current listing. Epsilon will use the original file pattern to rebuild the file listing. If you've marked files for copying, moving, or deleting, the markings will be discarded if you refresh, so Epsilon will prompt first to confirm that you want to do this.

The `S` key controls sorting. It prompts you to enter another letter to change the sorting method. Press `N`, `E`, `S`, or `D` to select sorting by file name, file extension, size, or time and date of modification, respectively. Press `U` to turn off sorting the next time Epsilon makes a dired listing, and display the file names in the same order they come from the operating system. (You can have Epsilon rebuild the current listing using the `R` subcommand.)

Press `+` or `-` at the sorting prompt to sort in ascending or descending order, respectively, or `R` to reverse the current sorting order.

Press `G` at the sorting prompt to toggle directory grouping. With directory grouping, Epsilon puts all subdirectories first in the list, then all files, and sorts each part individually. Without directory grouping, it mixes the two together (although it still puts `.` and `..` first).

Under Windows, press `A` to display the file's current attributes (Hidden, System, Read-only and Archive) and specify a new attribute list. You can set the `dired-layout` variable under Windows to include these attributes in the dired listing itself, or customize dired's format in other ways (for instance, to get wider listings with full dates including seconds for local Unix files). Under Unix, `A` runs the `chmod` command, passing it the mode specification you type, such as `g+w` to let group members write to the file. For remote files accessed via `Scp`, Epsilon sends the mode specification you provide directly to the `Sftp` server. It must be in the form of Unix-style octal permission bits, like `0644`.

Press `@` to toggle the read-only property of the current file. Under macOS, Linux and FreeBSD, Epsilon uses your `umask` when making a file writable, like `"chmod +w"`.

Press `V` to run the "viewer" for that file; the program assigned to it according to file associations. For programs, this will run the program. For document files, it typically opens the file in the program assigned to that file extension. For folders, it opens the folder in Windows File Explorer, macOS's Finder, or similar. See page 146 for information on associating Epsilon for Windows with particular file extensions.

Under Unix, `V` uses KDE, Gnome, or macOS file associations to run the viewer for the file. See the `epsilon-viewer` script to change which of these types of viewers Epsilon uses. For Gnome, run the `gnomecc` program to select a different viewer for a specific file type.

Press `Shift-P` to print the current file. Under Windows, press `T` to display the properties of a file or directory. (This is a convenient way to see the total size of all files in a directory.) Press `F` to search for text in a file name, skipping over matches in the columns for file size or date, by running incremental-search with a column restriction.

Several keys provide shortcuts for common operations. The `1` key examines the selected file in a window that occupies the whole screen (like typing `Ctrl-x 1 E`). The `2` key splits the current window horizontally and examines the selected file in the second window, leaving the dired buffer in the first (like typing `Ctrl-x 2 E`). The `5` key functions like the `2` key, but splits the window vertically (like typing `Ctrl-x 5 E`). The `0` key examines the selected file in the next window on the screen, without splitting windows any further. The `Z` key zooms the window to full-screen, then examines the selected file (like typing `Ctrl-x z e`).

Press `Shift-E` to examine the current file or directory, like `<Enter>`, but deleting the current dired buffer if you've moved to a new one. This runs the `dired-examine-deleting` function, while plain `E` runs `dired-examine`. You can swap these commands so plain `E` deletes old dired buffers while `Shift-E` doesn't, by adding these lines to your `init.el` customization file (see page 164):

```
~dired-tab "e": dired-examine-deleting
~dired-tab "E": dired-examine
```

(Similar lines can attach `dired-examine-deleting` to keys like `<Space>` or `<Enter>`. See page 158.)

Press lowercase `L` to create a live link. First Epsilon creates a second window, if there's only one window to start with. (Provide a numeric argument to get vertical, not horizontal, window splitting.) Then Epsilon displays the file named on the current dired line in that window, in a special live link buffer. As you move around in the dired buffer, the live link buffer will automatically update to display the current file. Files over `dired-live-link-limit` bytes in size won't be shown, to avoid delays. See the `wrap-dired-live-link` variable to control how long lines display. Delete the live link buffer or window, or show a different buffer there, to stop the live linking.

Press `Shift-G` to mark files by content. This subcommand prompts for some search text. You can use the keys `Ctrl-t`, `Ctrl-w` and `Ctrl-c` when typing the search string to toggle regex mode, word mode, and case folding.

Then the subcommand prompts for a key to indicate what kind of marking to apply. Press `d`, `m`, or `c` to mark files for deletion, moving or copying, `u` to remove markings, `U`, `L`, or `R` to perform the corresponding renaming function described above, or `g` to apply a generic marking that simply indicates which files contained a match for the search string. A numeric prefix argument to this subcommand reverses the sense of its test, marking only files that don't contain the specified text.

`Alt-[` and `Alt-]` move back and forward, respectively, by marks. They look at the mark on the current line (such as a `D` for deletion), then go to the next (or previous) line that has different markings. The `copy-file-name` command on `Ctrl-c Alt-n` copies the full pathname of the current line's file to the clipboard (just as it copies the current file's full pathname, in non-dired buffers).

Finally, typing `H` or `?` while in dired displays help on these dired subcommands.

4.8.9 Buffer List Editing

The `bufed` command on `Ctrl-x Ctrl-b` functions like dired, but it works with buffers instead of files. It creates a list of buffer names. Each buffer name appears on a line along with the size of the buffer, the associated file name (if any) and a star if the buffer contains unsaved changes, and/or an `R` if the buffer is currently marked read-only. The `bufed` command pops up the list, and highlights the line describing the current buffer.

In `bufed`'s popup window, alphabetic keys run special `bufed` commands. The `N` and `P` keys go to the next and previous buffers in the list, respectively, by going down or up one line. The `D` command deletes the buffer on the current line, but warns you if the buffer contains unsaved changes. The `S` key saves the buffer on the current line, and `Shift-P` prints the buffer like the `print-buffer` command. The `E` or `<Space>` command selects the buffer on the current line and displays it in the current window, removing the `bufed` listing.

As in dired, several keys provide shortcuts for common operations. The `1` key expands the current window to take up the whole screen, then selects the highlighted buffer. The `2` key splits the current window horizontally and selects the highlighted buffer in the second window. The `5` key works like the `2` key, except it splits the window vertically. The `Z` key zooms the current window to full-screen, then selects the highlighted buffer.

By default, the most recently accessed buffers appear at the top of the list, and those you haven't used recently appear at the end. The current buffer always appears at the top of the list. You can press `'b'`, `'f'`, or `'i'` to make Epsilon sort the list by buffer name, file name, or size, respectively. Pressing `'a'` makes Epsilon sort by access time again. Pressing the upper case letters `'B'`, `'F'`, `'I'`, or `'A'` reverses the sense of the sort. Pressing `'u'` produces a buffer list ordered by time of creation, with the oldest buffers at the bottom. Pressing `'m'` toggles whether modified buffers appear first in the list. Alphabetic keys not mentioned above do nothing. Most other keys like `<Down>` or `Ctrl-s` run their usual commands.

The `bufed` command does not normally list special buffers such as the kill buffers, whose names start with a dash character (`"-"`). To include even these buffers, give the `bufed` command a numeric argument.

You can set the `session-restore-directory` variable to control whether Epsilon restores any current directory setting in the session file. Set it to 0 and Epsilon will never do this. Set it to 1 and Epsilon will always restore the current directory when it reads a session file. The default value 2 makes Epsilon restore the current directory setting only when the `-w1` flag has been specified. (Under Windows, Epsilon's installer includes this flag when it makes Start Menu shortcuts.)

You can set the `session-restore-files` variable to control whether Epsilon restores files named in a session file, or just search strings, command history, and similar settings. If `session-restore-files` is 0, when Epsilon restores a session, it won't load any files named in the session, only things like previous search strings. If 1, the default, Epsilon will restore previous files as well as other settings. If 2, Epsilon will restore previous files only if there were no files specified on Epsilon's command line. The `session-always-restore` variable is more drastic, turning off session files entirely when there's a file specified on Epsilon's command line.

The `write-session` command writes a session file, detailing the files you're currently editing, the window configuration, default search strings, and so forth. By default, Epsilon writes a session file automatically whenever you exit, but you can use this command if you prefer to save and restore sessions manually.

The `read-session` command loads a session file, first asking if you want to save any unsaved files. Reading in a session file rereads any files mentioned in the session file, as well as replacing search strings, all bookmarks, and the window configuration. However, any files not mentioned in the session file will remain, as will keyboard macros, key bindings, and most variable settings. If you use either command and specify a different session file than the default, Epsilon will use the file name you provided when it automatically writes a session file as you exit.

Locating The Session File

By default, Epsilon restores your previous session by consulting a single session file named `epsilon.ses`, which is normally stored in your customization directory. (See page 13.) Epsilon will write such a file when you exit.

You can set Epsilon to use multiple session files by having it search for an existing session file, starting from the current directory. If a session file doesn't exist in the current directory, then Epsilon looks in its parent directory, then in that directory's parent, and so forth, until it reaches the root directory or finds a session file. Or you can have it always read and create its session file in the current directory.

To make Epsilon look for its session file only in the current directory, and create a new session file there on exiting, set the `session-default-directory` variable to `."`.

To make Epsilon search through a directory hierarchy for an existing session file, set the `session-tree-root` variable to empty. If this variable is set to a directory name in absolute form, Epsilon will only search for an existing session file in the named directory or one of its children. For example, if `session-tree-root` holds `c:\joe\proj`, and the current directory is `c:\joe\proj\src`, Epsilon will search in `c:\joe\proj\src`, then `c:\joe\proj`, for a session file. If the current directory is `c:\joe\misc`, on the other hand, Epsilon won't search at all (since `\joe\misc` isn't a child of `\joe\proj`), but will use the rules below. By default this variable is set to the word `NONE`, an impossible absolute directory name, so searching is disabled.

If Epsilon finds no such file by searching as described above (or if such searching is disabled, as it usually is), then Epsilon looks for a session file in each of these places, in this order:

- If the `session-default-directory` variable is non-empty, in the directory it names. (This variable is empty by default.)
- If the configuration variable `EPSPATH` can be found, in the first directory it names. (See page 10 for more on configuration variables.)

- In your customization directory.

There are three ways to tell Epsilon to search for a file with a different name, instead of the default of `epsilon.ses`. With any of these methods, specifying an absolute path keeps Epsilon from searching and forces it to use a particular file. Epsilon checks for alternate names in this order:

- The `-p` flag can specify a different session file name.
- An `ESESSION` configuration variable can specify a different session file name.
- The `session-file-name` variable can specify a name.

Summary: read-session
write-session

4.9.2 File Associations

You can set up file associations in Epsilon for Windows using the `configure-epsilon` command. It lets you modify a list of common extensions, then sets up Windows so if you double-click on a file with that extension, it invokes Epsilon to edit the file. The files will be sent to an existing copy of Epsilon, if one is running, or you can choose to always start a new instance.

To open a file in another program using file associations, you can use the `v` subcommand in `dired`, or `Alt-o v` (quick-`dired-command`) in any file.

Summary: configure-epsilon

4.9.3 Sending Files to a Prior Instance

Epsilon's command line flag `-add` tells Epsilon to locate an existing instance of itself (a "server"), send it a message containing the rest of the command line, and immediately exit. (Epsilon ignores the flag if there's no prior instance.)

The command line flag `-noserver` tells Epsilon that it should not respond to such messages from future instances.

The command line flag `-server` may be used to alter the server name for an instance of Epsilon, which is "Epsilon" by default. An instance of Epsilon started with `-server:somename` `-add` will only pass its command line to a previous instance started with the same `-server:somename` flag.

An `-add` message to Epsilon uses a subset of the syntax of Epsilon's command line. It can contain file names to edit, the `+linenum` flag, the `-dir dirname` flag to set a directory for interpreting any relative file names that follow, the flag `-dvarname=value` to set an Epsilon variable, `-lfilename` to load an EEL bytecode file, or `-rfuncname` to run an EEL function, command, or macro. Use `-rfuncname=arg` to run an EEL function and pass it a single string parameter *arg*. Epsilon unminimizes and tries to move to the foreground whenever it gets a message, unless the message uses the `-r` flag and doesn't include a file name.

Spaces separate file names and flags in the message; surround a file name or flag with " characters if it contains spaces. In EEL, such messages arrive via a special kind of `WIN_DRAG_DROP` event.

You can also use the `-wait` flag instead of `-add`. This causes the client Epsilon to send the following command line to an existing instance and then wait for a response from the server, indicating the user has finished editing the specified file. Use the `resume-client` command on `Ctrl-c #` to indicate this.

Epsilon for Windows normally acts as a server for its own internal-format messages, as described above, and also acts as a DDE server for messages from Windows File Explorer. The `-noserver` flag described above also disables DDE, and the `-server` flag also sets the DDE server name. The DDE server in Epsilon uses a topic name of “Open” and a server name determined as described above (normally “Epsilon”).

When Epsilon gets an `-add` message, it tries to move itself to the top of the window order (unless the message used the `-r` flag and specified no file; then the function run via `-r` is responsible for changing the window order, if desired). Under Epsilon for X11, the `server-raises-window` variable controls this behavior.

Summary: Ctrl-c # resume-client

4.9.4 MS-Windows Integration Features

Epsilon can integrate with Microsoft’s Visual Studio (Developer Studio) in several ways. The on-demand style of integration lets you press a key (or click a button) while editing a file in Visual Studio, and start Epsilon on the same file. The other automates this process, so any attempt to open a source file in Visual Studio is routed to Epsilon.

For on-demand integration, you can add Epsilon to the Tools menu in Microsoft Visual Studio. You’ll then be able to select Epsilon from the menu and have it begin editing the same file you’re viewing in Visual Studio, at the same line.

To do this in Visual Studio 5.0 or 6.0, use the Tools/Customize menu command in Visual Studio. Select the Tools tab in the Customize dialog that appears. Create a new entry for the Tools menu, and set the Command field to the name of Epsilon’s executable, `epsilon.exe`. Include its full path, typically `c:\Program Files\Eps14\Bin\epsilon.exe`. Set the Arguments field to `-add +$(CurLine):$(CurCol) $(FilePath)`. Set the Initial Directory field to `$(FileDir)`. After creating this new command, you can then use the Tools/Customize/Keyboard command to set up a shortcut key for it. You may also want to configure Visual Studio to detect when a file is changed outside its environment, and automatically load it. See Tools/Options/Editor for this setting.

If you use Visual Studio .NET, the steps are slightly different. Use Tools/External Tools/Add to create a new entry in the Tools menu. Set the Command field to the full path to Epsilon’s executable, `epsilon.exe`, as above. Set the Arguments field to `-add +$(CurLine):$(CurCol) $(ItemPath)`. Optionally, set the Initial Directory field to `$(ItemDir)`. You can use Tools/Options/Environment/Keyboard to set up a shortcut key for the appropriate Tools.ExternalCommand entry. To set Visual Studio .NET to autoload modified files, use Tools/Options/Environment/Documents.

You can also set up Visual Studio 5.0 or 6.0 so that every time Visual Studio tries to open a source file, Epsilon appears and opens the file instead. To set up Visual Studio so its attempts to open a source file are passed to Epsilon, use the Customize command on the Tools menu and select the Add-ins and Macro Files page in the dialog. Click Browse, select Add-ins (.dll) as the File Type, and navigate to the `VisEpsil.dll` file located in the directory containing Epsilon’s executable (typically `c:\Program Files\Eps14\bin`). Select that file.

Close the Customize dialog and a window containing an Epsilon icon (a blue letter E) should appear. You can move the icon to any toolbar by dragging it. Click the icon and a dialog will appear with two options. Unchecking the first will disable this add-in entirely. If you uncheck the second, then any time you try to open a text file in Dev Studio it will open in both Epsilon and Dev Studio. When checked, it will only open in Epsilon.

Running Epsilon via a Shortcut

Epsilon comes with a program, `sendeps.exe`, that's installed in the directory containing Epsilon's main executable. It provides some flexibility when you create a desktop icon for Epsilon, or use the Send To feature (both of which involve creating a Windows shortcut).

When Epsilon's Windows installer or its `configure-epsilon` command create a desktop shortcut for Epsilon, or an entry on Explorer's Send To menu, they're set to run this `sendeps.exe` program instead of Epsilon's main executable. Sendeps will start Epsilon if necessary, or locate an existing copy of Epsilon, and load the files named on its command line.

This is useful because Windows ignores a shortcut's flags (command line settings) when you drop a document on a shortcut, or when you use the Send To feature. (If it used the flags, you could simply create a shortcut to Epsilon's main executable and pass its `-add` flag. Since it doesn't, sending a file requires a separate program.) Also, Windows sends long file names without quoting them in these cases, which would cause problems if sent directly to Epsilon. (Don't use `sendeps.exe` for pinned taskbar entries, since those don't have such issues, and will result in Windows showing an extra Epsilon taskbar icon. Pin Epsilon to the taskbar by right-clicking its taskbar icon.)

`Sendeps.exe` is configured through entries in a `lugeps.ini` file located in `sendeps's` directory. Normally you would run the `configure-epsilon` command to modify this file, but we document its format here for completeness and for special uses.

The `lugeps.ini` file may have multiple sections, each section name inside square brackets. The section name `sendeps` uses is taken from the base name of its executable. That means making copies of the `sendeps.exe` program under different names lets you have multiple Send To or desktop entries that behave differently. The `configure-epsilon` command makes a copy of `sendeps.exe` named `sendnew.exe`, for example, to support certain configuration options.

These are the default settings for `lugeps.ini`:

```
[SendEps]
server=Epsilon
topic=Open
ddeflags=
executable=epsilon.exe
runflags=-add -w1
nofilestartnew=1
nofileflags=-w1
usedde=0
senddir=1
```

Here's how Sendeps uses the above settings. It first looks for an Epsilon server named `server` using Epsilon's `-add` protocol. If found, it sends the server a command line consisting of the `ddeflags` setting, followed by the file name passed on its command line (inside double quotes). If there's no such server running, Sendeps executes a command line built by concatenating the `executable` name, the `runflags`, and the quoted file name. If the executable name is a relative pathname, Epsilon searches for it first in the directory containing `sendeps.exe`, then along your `PATH` environment variable.

Normally relative file names on the `sendeps` command line are sent to Epsilon as-is, along with a `-dir` flag indicating Sendeps's current directory. Set `senddir` to zero and Sendeps will convert each file name to absolute form itself and omit `-dir`.

You can tell Sendeps to use DDE instead of its usual `-add` protocol by setting `usedde` to 1. In that case it will use the specified `topic` name.

When you invoke Sendeps without specifying a file name on its command line, its behavior is controlled by the `nofilestartnew` setting. If nonzero, it starts a new instance of Epsilon. If zero, it brings an existing instance to the top, if there is one, and starts a new instance otherwise. In either case, if it needs to start a new instance, it uses `nofileflags` on the command line.

The Open With Epsilon Shell Extension

Epsilon includes a 32-bit shell extension DLL named `owitheps.dll`, which can be used to add an Open With Epsilon entry to every file's context menu in 32-bit versions of File Explorer. But in this version of Epsilon, the `configure-epsilon` accomplishes this by simply using registry entries (since this works in all Windows versions), so this shell extension is no longer used. We document it here in case you have an older Epsilon installation that uses the extension.

The Open With Epsilon shell extension may be configured by creating entries in a file named `lugeps.ini` located in your Windows directory.

These are its default settings, which you can copy to `lugeps.ini` as a basis for your changes:

```
[OpenWith]
server=Epsilon
serverflags=
executable=epsilon.exe
runflags=-add -w1
menutext=Open With Epsilon
```

Once installed, if you select Open With Epsilon from the menu in File Explorer, the shell extension first looks for an Epsilon server named `server` using Epsilon's `-add` protocol. If found, it sends the server a command line consisting of the `serverflags` setting, followed by the file name you selected (inside double quotes).

If there's no such server running, the DLL executes a command line built by concatenating the `executable` name, the `runflags`, and the quoted file name. If the `executable` name is a relative pathname, it first tries to run any executable by that name located in the DLL's current directory. If that fails, it uses the `executable` name as-is, and lets Windows search for it along the `PATH`.

If you've selected multiple files, it repeats the above process for each file.

You can alter the menu text Explorer displays by setting the `menutext` item. This setting doesn't take effect until you restart Explorer, or unload and reload the `owitheps.dll` file that provides this menu by running `regsvr32 /u owitheps.dll`, then `regsvr32 owitheps.dll`. Other changes to the DLL's settings take effect immediately.

4.10 Running Other Programs

Epsilon provides several methods for running other programs from within Epsilon. The `push` command on `Ctrl-x Ctrl-e` starts a command processor (shell) running. You can then issue shell commands. When you type the "exit" command, you will return to Epsilon and can resume your work right where you left off.

With a numeric argument, the command asks for a command line to pass to the shell, runs this command, then returns.

While Epsilon runs a command processor or other program with the `push` command, it looks like you ran the program from outside of Epsilon. But Epsilon can make a copy of the input and output that occurs during the program's execution, and show it to you when the program returns to Epsilon. If you set the

variable `capture-output` to a nonzero value (normally it has the value zero), Epsilon will make such a transcript. When you return to Epsilon, this transcript will appear in a buffer named “process”.

You can use the `filter-region` command on Alt-I to process the current region through an external command. Epsilon will run the command, sending a copy of the region to it as its standard input. By default, the external command’s output goes to a new buffer named “filter-output”, but with a negative prefix argument the command asks for the buffer name to use for output. Run `filter-region` with a non-negative numeric argument if you want the output to replace the current region.

The command `shell-command` is similar, but it doesn’t send the current region as the command’s input. It prompts for the name of an external program to run and displays the result in a buffer; with a non-negative numeric argument it inserts the command’s output into the current buffer and with a negative prefix argument it prompts for the buffer to receive the output.

Configuration variables (see page 10) let you customize what command Epsilon runs when it wants to start a process. Epsilon runs the command file named by the `EPSCOMSPEC` configuration variable. If no such variable exists, Epsilon uses the standard `COMSPEC` environment variable instead. Epsilon reports an error if neither exists.

If a configuration variable named `INTERSHELLFLAGS` has been defined, Epsilon passes the contents of this variable to the program as its command line. When Epsilon needs to pass a command line to the program, it doesn’t use `INTERSHELLFLAGS`. Instead, it inserts the contents of the `CMDSHELLFLAGS` variable before the command line you type.

The sequence `%%` in `CMDSHELLFLAGS` makes Epsilon interpolate the command line at that point, instead of adding it after the flags. Put a “d” after the `%%` to have Epsilon double backslashes in the command line until the first space; put “b” to have Epsilon change backslashes to slashes until the first space, or “a” to have Epsilon interpolate the command line as-is. A plain `%%` makes Epsilon guess: it uses “b” if the shell name ends in “sh”, otherwise “a”.

If Epsilon can’t find a definition for `INTERSHELLFLAGS` or `CMDSHELLFLAGS`, it substitutes flags appropriate for the operating system. See the next section for more on these settings.

Summary:	Ctrl-x Ctrl-e	push
		filter-region
		shell-command

4.10.1 The Concurrent Process

Epsilon can also run a program in a special way that allows you to interact with the program in a buffer and continue editing while the program runs. It can help in preparing command lines, by letting you edit things you previously typed, and it automatically saves what each program types, so you can examine it later. If a program takes a long time to produce a result, you can continue to edit files while it works. We call a program run in this way a *concurrent process*.

The `start-process` command, bound to Ctrl-x Ctrl-m, begins a concurrent process. Without a numeric argument, it starts a shell command processor which will run until you exit it (by going to the end of the buffer and typing “exit”). With a numeric argument, it creates an additional process buffer, in Epsilon environments that support more than one.

Epsilon maintains a command history for the concurrent process buffer. You can use Alt-p and Alt-n to retrieve the text of previous commands. With a numeric prefix argument, these keys show a menu of all previous commands. You can select one to repeat.

In a concurrent process buffer, you can use the `<Tab>` key to perform completion on file names, command names, and environments variable names as you're typing them. If no more completion is possible, Epsilon displays all the matches in the echo area, if they fit. If not, press `<Tab>` again to see them listed in the buffer. See the `process-complete` command for more details, or the `process-completion-style` and `process-completion-dircmds` variables to customize how process buffer completion works.

The `process-coloring-rules` variable controls how Epsilon interprets ANSI escape sequences and similar in a process buffer, and `process-echo` changes certain echoing functions. The `process-enter-whole-line` variable customizes how the process buffer behaves when you press `<Enter>`, or select a previous command from command history.

Under Windows, certain process buffer functions like completion, or interpreting compiler error messages, require Epsilon to determine the current directory of the command processor running there. Epsilon does this by examining each prompt from `cmd.exe`, such as `C:\WINNT>`. If you've set a different format for the prompt, you may have to set the `process-prompt-pattern` variable to tell Epsilon how to retrieve the directory name from it. Also see the `use-process-current-directory` variable to change how Epsilon's current directory and the process's are linked together.

As described in the previous section, you can change the name of the shell command processor Epsilon calls, and specify what command line switches Epsilon should pass to it, by setting configuration variables. Some different configuration variable names override those variables, but only when Epsilon starts a subprocess concurrently. For example, you might run a command processor that you have to start with a special flag when Epsilon runs it concurrently. The `INTERCONCURSHELLFLAGS` and `CMDCONCURSHELLFLAGS` variables override `INTERSHELLFLAGS` and `CMDSHELLFLAGS`, respectively. The `EPSCONCURCOMSPEC` variable overrides `EPSCOMSPEC`.

For example, one version of the Bash shell for Windows systems requires these settings:

```
EpsComspec=c:\cygwin\bin\bash.exe
InterShellFlags=--login --noediting -i
CmdShellFlags=-c "%%"
```

These are configuration variables, so they go in the environment for Epsilon for Unix, or in the system registry for Windows versions. See page 10. With some replacement shells, you may also have to set the `process-echo` variable. (Cygwin users: also see the `cygwin-filenames` variable.)

When a concurrent process starts, Epsilon creates a buffer named "process". In this buffer, you can see what the process types and respond to the process's requests for input. If a buffer named "process" already exists, perhaps from running a process previously, Epsilon goes to its end. Provide a numeric argument to the `start-process` command and it will create an additional process buffer (in those environments where Epsilon supports multiple process buffers).

If you set the variable `clear-process-buffer` to 1, the commands `start-process`, `push`, and `make` (described below) will each begin by emptying the process buffer. The variable normally has a value of 0. (See that variable for more options.) Set the variable `start-process-in-buffer-directory` to control which directory the new process starts in.

A program running concurrently behaves as it does when run directly from outside Epsilon except when it prints things on the screen or reads characters from the keyboard. When the program prints characters, Epsilon inserts these in the process buffer. When the program waits for a line of input, Epsilon will suspend the process until it can read a line of input from the process buffer, at which time Epsilon will restart the process and give it the line of input. You can type lines of input before the program requests them, and Epsilon will feed the input to the process as it requests each line. In some environments, Epsilon will also satisfy requests from the concurrent process for single-character input.

In detail, Epsilon remembers a particular spot in the process buffer where all input and output takes place. This spot, called the *type point*, determines what characters from the buffer a program will read when it does input, and where the characters a program types will appear. Epsilon inserts in the buffer, just before the type point, each character a program types. When a process requests a line of input, Epsilon waits until a newline appears in the buffer after the type point, then gives the line to the program, then moves the type point past these characters. (In environments where Epsilon can distinguish a request by a program to read a single character, it will pause the concurrent process until you have inserted a character after the type point, give that character to the concurrent process, then advance the type point past that character.)

You may insert characters into the process buffer in any way you please, typing them directly or using the yank command to retrieve program input from somewhere else. (Also see the `process-yank-confirm` variable.) You can move about in the process buffer, edit other files, or do anything else at any time, regardless of whether the program has asked the system for keyboard input.

To generate an end-of-file condition for DOS or Windows programs reading from the standard input, insert a `^Z` character by typing `Ctrl-q Ctrl-z` on a line by itself, at the end of the buffer. For a Unix program, type `Ctrl-q Ctrl-d` `<Enter>`.

Some programs will not work when running concurrently. Programs that do cursor positioning or graphics will not work well, since such things do not correspond to a stream of characters coming from the program to insert into a buffer. They may even interfere with what Epsilon displays. We provide the concurrent process facility primarily to let you run programs like compilers, linkers, assemblers, filters, etc.

You may rename the buffer named “process” using the `rename-buffer` command, and start a different, independent concurrent process in the buffer “process”. Or run the `start-process` command with a numeric argument to have Epsilon pick a unique buffer name for the new process buffer if “process” already has an active process. If you exit Epsilon while running a concurrent process, Epsilon kills that process.

The `exit-process` command types “exit” to a running concurrent process. If the concurrent process is running a standard command processor, it should then exit. Also see the `process-warn-on-exit` and `process-warn-on-killing` variables.

The `kill-process` command disconnects Epsilon from a concurrent process, and forces it to exit. It operates on the current buffer’s process, if any, or on the buffer named “process” if the current buffer has no process. If the current buffer isn’t a process buffer but has a running Internet job (such as an `ftp://` buffer), this command tries to cancel it.

The `stop-process` command, normally on `Ctrl-c Ctrl-c`, makes a program running concurrently believe you typed Control-Break (or, for Unix, sends an interrupt signal). It operates on the current buffer’s process, if any, or on the buffer named “process” if the current buffer has no process.

Summary:	<code>Ctrl-x Ctrl-m</code>	<code>start-process</code>
	<code>Ctrl-c Ctrl-c</code>	<code>stop-process</code>
	Process mode only: <code>Alt-⟨Backspace⟩</code>	<code>process-backward-kill-word</code>
	Process mode only: <code>⟨Tab⟩</code>	<code>process-complete</code>
	Process mode only: <code>C-Y</code>	<code>process-yank</code>
	Process mode only: <code>Alt-n</code>	<code>process-next-cmd</code>
	Process mode only: <code>Alt-p</code>	<code>process-previous-cmd</code>
		<code>kill-process</code>
		<code>exit-process</code>

4.10.2 Compiling From Epsilon

Many compilers produce error messages in a format that Epsilon can interpret with its `next-error` command on `Ctrl-x Ctrl-n`. The command searches in the process buffer (beginning at the place it reached last time, or at the beginning of the last command) for a line that contains a file name, a line number, and an error message. If it finds one, it uses the `find-file` command to retrieve the file (if not already in a window), then goes to the appropriate line in the file. With a numeric argument, it finds the n th next error message, or the n th previous one if negative. In particular, a numeric argument of 0 repeats the last message. The `previous-error` command on `Ctrl-x Ctrl-p` works similarly, except that it searches backward instead of forward.

The `Ctrl-x Ctrl-n` and `Ctrl-x Ctrl-p` keys move back and forth over the list of errors. If you move point around in a process buffer, it doesn't change the current error message. You can use the `find-linked-file` command on `Ctrl-x Ctrl-l` to reset the current error message to the one shown on the current line. (The command also goes to the indicated source file and line, like `Ctrl-x Ctrl-n` would.)

Actually, `Ctrl-x Ctrl-n` runs the `next-position` command, not `next-error`. The `next-position` command usually calls `next-error`. After you use the `grep` command (see page 48), however, `next-position` calls `next-match` instead, to move to the next match of the pattern you searched for. If you use any command that runs a process, or run `next-error` explicitly, then `next-position` will again call `next-error` to move to the next error message.

Similarly, `Ctrl-x Ctrl-p` actually runs `previous-position`, which decides whether to call `previous-error` or `previous-match` based on whether you last ran a compiler or searched across files.

To locate error messages, the `next-error` command performs a regular-expression search using a pattern that matches most compiler error messages. See page 66 for an explanation of regular expressions. The command uses the `ERROR_PATTERN` macro and others, defined in the file `nexterr.e`. You can change these patterns if they don't match your compiler's error message format. The `next-error` command also uses another regular-expression pattern to filter out any error messages Epsilon should skip over, even if they match `ERROR_PATTERN`. The variable `ignore-error` stores this regular expression. For example, if `ignore-error` contains the pattern `".*warning"`, Epsilon will skip over any error messages that contain the word "warning".

The `next-error` command knows how to use the Java `CLASSPATH` to locate Java source files, and has special logic for running Cygwin-based programs under Windows. (See the `cygwin-filenames` variable for more information on the latter.) You can set the `process-next-error-options` variable to control how this command looks for a file.

If you run a compiler via telnet or a similar process in an Epsilon buffer, you can set up `next-error` to translate file names in the telnet buffer into URL-style file names that Epsilon can use to access the file. See page 132.

The command `view-process` on `Shift-F3` can be convenient when there are many long error messages in a compilation. It pops up a window showing the process buffer and its error messages, and lets you move to a particular line with an error message and press `<Enter>`. It then goes to the source file and line in error. You can also use it to see the complete error message from the compiler, when `next-error`'s one-line display is inadequate.

An alternative to using `view-process` is setting the `process-view-error-lines` variable nonzero. It tells the `next-error` and `previous-error` commands to ensure the error in the process buffer is visible in a window each time it moves to a source line.

The `make` command on `Ctrl-x m` runs a program and scans its output for error messages using `next-error`. In some environments, it runs the program in the background, displaying a "Compiling" message while it runs. If you don't start other editing before it finishes, it automatically goes to the first error message. If you do, it skips this step; you can press `Ctrl-x Ctrl-n` as usual to go to the first error. In other

environments, Epsilon may run the program in a concurrent process buffer, or non-concurrently. See the variables `compile-in-separate-buffer` and `concurrent-make` for more details.

By default, it runs a program called “make”, but with a numeric argument it will prompt for the command line to execute. It will use that command line from then on, if you invoke `make` without a numeric argument. See the variable `start-make-in-buffer-directory` to control which directory the new process starts in.

Epsilon uses a template for the command line (stored in the `push-cmd` variable), so you can define a command line that depends on the current file name. See page 122 for information on templates. For example, `c1 %f` runs the `c1` command, passing it the current file name.

If a concurrent process already exists, Epsilon will attempt to run the program concurrently by typing its name at the end of the process buffer (in those environments where Epsilon isn’t capable of creating more than one process buffer). When Epsilon uses an existing process buffer in this way, it will run `next-error` only if you’ve typed no keys during the execution of the concurrent program. You can set the variable `concurrent-make` to 0 to force Epsilon to exit any concurrent process, before running the “make” command. Set it to 2 to force Epsilon to run the command concurrently, starting a new concurrent process if it needs to. When the variable is 1 (the default), the `make` command runs the compiler concurrently if a concurrent process is already running, non-concurrently otherwise.

Whenever `push` or `make` exit from a concurrent process to run a command non-concurrently, they will restart the concurrent process once the command finishes. Set the `restart-concurrent` variable to zero if you don’t want Epsilon to restart the concurrent process in this case.

Before `make` runs the program, it checks to see if you have any unsaved buffers. If you do, it asks if it should save them first, displaying the buffers using the `bufed` command. If you say yes, then the `make` command saves all of your unsaved buffers using the `save-all-buffers` command (which you can also invoke yourself with `Ctrl-x s`). You can modify the `save-when-making` variable to change this behavior. If it has a value of 0, Epsilon won’t warn you that you have unsaved buffers. If it has a value of 1, Epsilon will automatically save all the buffers without asking. If it has a value of 2 (as it has normally), Epsilon asks.

The `compile-buffer` command on `Alt-F3` is somewhat similar to `make`, but tries to compile only the current file, based on its extension. There are several variables like `compile-cpp-cmd` you can set to tell Epsilon the appropriate compilation command for each extension. If Epsilon doesn’t know how to compile a certain type of file, it will prompt for a command line. While Epsilon’s `make` command is good for compiling entire projects, `compile-buffer` is handy for compiling simple, one-file programs.

The command is especially convenient for EEL programmers because `compile-buffer` automatically loads the EEL program into Epsilon after compiling it. The EEL compiler is integrated into Epsilon, so Epsilon doesn’t need to run another program to compile. When Epsilon compiles EEL code using its internal EEL compiler, it looks in the `compile-eel-dll-flags` variable for EEL command line flags.

The buffer-specific `concurrent-compile` variable tells `compile-buffer` whether to run the compiler concurrently. The value 2 means always run the compiler concurrently, 0 means never run concurrently, and 1 means run concurrently if and only if a concurrent process is already running. The value 3 (the default) means use the value of the variable `concurrent-make` instead. (The `concurrent-make` variable tells the `make` command whether to run its program concurrently, and takes on values of 0, 1, or 2 with the same meaning as for `concurrent-compile`.)

A file can use a file variable named “`compile-command`” (see page 127) to tell `compile-buffer` to use a specific command to compile that file, not the usual one implied by its file name extension. The command line must be surrounded with " characters if it contains a semicolon. For instance,

```
--* compile-command: "gcc -I/usr/local/include %r" --*
```

on the first line of a file tells Epsilon to use that command to compile that file. Unlike other file variables, Epsilon doesn’t scan for a compile command when you first load the file; it does this each time you use the

compile-buffer command. Also see the variable `use-compile-command-file-variable`.

Summary:	Ctrl-x Ctrl-n	next-position
	Ctrl-x Ctrl-p	previous-position
		next-error
		previous-error
	Shift-F3	view-process
	Ctrl-x m	make
	Alt-F3	compile-buffer

4.11 Repeating Commands

4.11.1 Repeating a Single Command

You may give any Epsilon command a numeric prefix argument. Numeric arguments can go up to several hundred million, and can have either a positive or negative sign. Epsilon commands, unless stated otherwise in their description, use a numeric argument as a repetition count if this makes sense. For instance, `forward-word` goes forward 10 words if given a numeric argument of 10, or goes backward 3 words if given a numeric argument of -3 .

The `argument` command, normally bound to `Ctrl-u`, specifies a numeric argument. After typing `Ctrl-u`, type a sequence of digits and then the command to which to apply the numeric argument. Typing a minus sign changes the sign of the numeric argument.

You may also use the Alt versions of the digit keys (Alt-1, etc.) with this command. (Note that by default the numeric keypad keys plus Alt do not give Alt digits. They produce keys like Alt-⟨PgUp⟩ or let you enter special characters by their numeric code. See the `alt-numeric-keys` variable.) You can enter a numeric argument by holding down the Alt key and typing the number on the main keyboard. Alt-⟨Minus⟩ will change the sign of a numeric argument, or start one at -4 .

If you omit the digits, and just say `Ctrl-u Ctrl-f`, for instance, Epsilon will provide a default numeric argument of 4 and move forward four characters. Typing another `Ctrl-u` after invoking `argument` multiplies the current numeric argument by four, so typing `Ctrl-u Ctrl-u Ctrl-n` will move down sixteen lines. In general typing a sequence of n `Ctrl-u`'s will produce a numeric argument of 4^n .

The `run-with-argument` command provides an alternative way to run a command with a numeric argument. It prompts for the argument with a normal Epsilon numeric prompt, so that you can yank the number to use from the clipboard, or specify it with a different base like `0x1000` (for hexadecimal), or specify the number as a character code like `'q'` or `'\n'` or `<Yen Sign>`.

Summary:	Ctrl-u	argument
		run-with-argument

4.11.2 Keyboard Macros

Epsilon can remember a set of keystrokes, and store them away in a *keyboard macro*. Executing a keyboard macro has the same effect as typing the characters themselves. Use keyboard macros to make repetitive changes to a buffer that involve the same keystrokes. You can even write new commands with keyboard macros.

Ctrl-x)	end-kbd-macro
Ctrl-F4, Ctrl-x e	last-kbd-macro
Shift-F4	pause-macro
Ctrl-x Alt-n	bind-last-macro
	name-kbd-macro
	insert-macro
	load-buffer
	redisplay

4.12 Simple Customizing

4.12.1 Bindings

Epsilon allows you to create your own commands and attach them, or any pre-existing Epsilon commands, to any key. If you bind a command to a key, you can then invoke that command by pressing the key. For example, at startup, Epsilon has `forward-character` bound to the Ctrl-f key. By typing Ctrl-f, the `forward-character` command executes, so point moves forward one character. If you prefer to have the command which moves point to the end of the current line, `end-of-line`, bound to Ctrl-f, you may bind that there.

You bind commands to keys with the `bind-to-key` command, which you can invoke with the F4 key. The `bind-to-key` command asks you for the name of a command (with completion), and the key to which to bind that command. You may precede the key by any number of *prefix keys*. When you type a prefix key, Epsilon asks you for another key. For example, if you type Ctrl-x, Epsilon asks you for another key. Suppose you type Ctrl-o. Epsilon would then bind the command to the Ctrl-x Ctrl-o key sequence. Prefix keys give Epsilon a virtually unlimited number of keys.

Epsilon at startup provides Ctrl-x and Ctrl-c as the only prefix keys. You can invoke many commands, such as `save-file` (Ctrl-x Ctrl-s) and `find-file` (Ctrl-x Ctrl-f), through the Ctrl-x prefix key. You may define your own prefix keys with the command called `create-prefix-command`. Epsilon asks you for a key to make into a prefix key. You may then bind commands to keys prefixed with this key using the `bind-to-key` command. To remove prefix keys, see page 166.

When you press a prefix key, Epsilon displays the key in the echo area to indicate that you must type another key. Epsilon normally displays the key immediately, but you can make it pause for a moment before displaying the key. If you press another key during the pause, Epsilon doesn't bother displaying the first key.

You control the amount of time Epsilon pauses using the `mention-delay` variable, expressed in tenths of a second. By default, this variable has a value of zero, which indicates no delay. You may find it useful to set `mention-delay` to a small value (perhaps 3). This delay applies in most situations where Epsilon prompts for a single key, such as when entering a numeric argument.

The `unbind-key` command asks for a key and then offers to rebind the key to the `normal-character` command, or to remove any binding it may have. A key bound to `normal-character` will self-insert; that's how keys like 'j' are bound. A key with no binding at all simply displays an error message.

You may bind a given command to any number of keys. You may invoke a command, whether or not bound to a key, using `named-command`, by pressing the Alt-x key. Alt-x asks for the name of a command, then runs the command you specified. This command passes any numeric argument you give it to the command it invokes. (Alt-x, like binding, also works with keyboard macros. And it works with any EEL subroutine that takes no parameters, even if it was not defined using the "command" keyword.)

Some keys behave differently in different contexts. For instance, the `n` key in `dire` mode moves down one line instead of inserting an “n” as it does in `Fundamental` mode. Every mode has a key table that defines which keys have “custom” bindings for that mode. Epsilon uses other key tables when you’re typing a response at a prompt, or scrolling through a list of choices. You can customize those bindings too.

The `bind-to-key` command modifies the current mode’s key binding if the selected key has a mode-specific binding in the current mode. If not, it modifies the global key table, so the binding applies to all modes.

To interactively create a new mode-specific binding, or change the bindings of keys in special modes like the ones used for prompts, the easiest method is to use the `list-all` and `load-buffer` commands. The former lists all the current bindings, so you can see the names of the internal commands run by each key and combine the lines to specific the bindings you want.

For example, at many prompts “?” lists possible responses. Say you want the `F8` key to do this too. In most contexts, `F8` runs the `set-variable` command. Examine the output of `list-all` and you’ll find lines like these:

```
~reg-tab "F-8": set-variable
...
~comp-tab "?": inp-show-matches
```

The first defines the usual binding for `F8`, showing how to specify that key in command line syntax, as `F-8`. The second specifies the binding for “?” at prompts, showing that it runs the internal function `inp-show-matches`. So combine them into

```
~comp-tab "F-8": inp-show-matches
```

Put the above line in a buffer and run `Alt-x load-buffer` on it, and now `F8` at prompts will display matches. Put that line in your `init.ecm` file (see page 162) and Epsilon will load the binding every time it starts up.

The command `alt-prefix`, bound to `<Esc>`, gets another key and executes the command bound to the `Alt` version of that key. You will find this command useful if you must use Epsilon from a keyboard lacking a working `Alt` key, or if you prefer to avoid using `Alt` keys. Also, you may find some combinations of control and `alt` awkward to type on some keyboards. For example, some people prefer to invoke the `replace-string` command by typing `<Esc> &` rather than by typing `Alt-&`.

The command `ctrl-prefix`, bound to `Ctrl-^`, functions similarly. It gets another key and converts it into the `Control` version of that key. For example, it changes ‘s’ into the `Ctrl-s` key.

Some key combinations are variations of other key combinations. For instance, `<Backspace>` and `Ctrl-h` are related in this way. Epsilon uses a notion of generic versus specific keys; for instance, the specific key `<Backspace>` is also generically a `Ctrl-h` key. If you bind this key to a new command, Epsilon will ask if you want to bind only the `<Backspace>` key, or all key combinations that generate a `Ctrl-h`.

Summary:	<code>Alt-x, F2</code>	<code>named-command</code>
	<code>F4</code>	<code>bind-to-key</code>
		<code>create-prefix-command</code>
		<code>unbind-key</code>
	<code><Esc></code>	<code>alt-prefix</code>
	<code>Ctrl-^</code>	<code>ctrl-prefix</code>

4.12.2 Brief Emulation

Epsilon can emulate the Brief text editor. The `brief-keyboard` command loads a Brief-style keyboard map. To undo this change, you can use the `epsilon-keyboard` command, which restores the standard keyboard configuration. This command only modifies those key combinations that Brief uses. Other keys retain their Epsilon definition. The Brief key map appears in figure 4.9.

In this release, Epsilon doesn't emulate a few parts of Brief. The separate command for toggling regular expression mode is not present, but you can type `Ctrl-t` within any searching command to toggle it. Regular expressions follow Epsilon's syntax, not Brief's. Brief's commands for loading and saving keyboard macro files aren't implemented, since Epsilon lets you have an unlimited number of macros loaded at once, not just one. Epsilon will beep if you press the key of an unimplemented Brief emulation command.

In Brief, the shifted arrow keys normally switch windows. But Epsilon adopts the Windows convention that shifted arrow keys select text. In Brief mode, the `Alt-arrow` keys on the separate cursor pad may be used to switch windows.

You can make Epsilon's display resemble Brief's display using the `set-display-look` command. See page 115.

4.12.3 CUA Keyboard

In CUA emulation mode, Epsilon recognizes most of the key combinations commonly used in Windows programs. Other keys generally retain their usual Epsilon function.

To enable this emulation, press `Alt-x`, then type `cua-keyboard` and press `(Enter)`. Use `Alt-x epsilon-keyboard (Enter)` to return to Epsilon's default key assignments.

The table shows the CUA key combinations that differ from Epsilon's native (Emacs-style) key configuration. In addition, various `Alt-letter` key combinations not mentioned here invoke menu items (for example, `Alt-f` displays the File menu in CUA mode, though it doesn't in Epsilon's native configuration).

Many commands in Epsilon are two-key combinations starting with `Ctrl-x` or `Ctrl-c`. In CUA mode, use `Ctrl-w` instead of `Ctrl-x`, and `Ctrl-k` instead of `Ctrl-c`. For example, the command `delete-blank-lines`, normally on `Ctrl-x Ctrl-o`, is on `Ctrl-w Ctrl-o` in CUA emulation.

4.12.4 Variables

You can set any user variable with the `set-variable` command. The variable must have the type *byte*, *char*, *short*, *int*, *array of chars*, or *pointer to char*. The command first asks you for the name of the variable to set. You can use completion. After you select the variable, the command asks you for the new value. Then the command shows you the new value.

Whenever Epsilon asks you for a number, as in the `set-variable` command, it normally interprets the number you give in base 10. But you can enter a number in hexadecimal (base 16) by beginning the number with `"0x"`, just like EEL integer constants. The prefix `"0o"` means octal, and `"0b"` means binary. For example, the numbers `"30"`, `"0x1E"`, `"0o36"`, and `"0b11110"` all refer to the same number, thirty. You can also specify an ASCII value by enclosing the character in single quotes. For example, you could type `'a'` to specify the ASCII value of the character `"a"` (in this example, 97). Or specify a Unicode character name inside `<` and `>` characters. You can also enter an arithmetic expression (anything the `eval` command can evaluate as a number).

The `set-any-variable` command is similar to `set-variable`, but also includes *system variables*. Epsilon uses system variables to implement its commands; unless you're writing EEL extensions, there's generally no reason to set them. When an EEL program defines a new variable, Epsilon considers it a system variable unless the definition includes the user keyword.

Alt-a	mark-normal-region	Alt-1	brief-drop-bookmark 1
Alt-b	bufed	Alt-2	brief-drop-bookmark 2
Ctrl-b	line-to-bottom
Ctrl-c	center-window	Alt-0	brief-drop-bookmark 10
Alt-c	mark-rectangle	F1	move-to-window
Alt-d	kill-current-line	Alt-F1	toggle-borders
Ctrl-d	scroll-down	F2	brief-resize-window
Alt-e	find-file	Alt-F2	zoom-window
Ctrl-e	scroll-up	F3	brief-split-window
Alt-f	display-buffer-info	F4	brief-delete-window
Alt-g	goto-line	F5	string-search
Alt-h	help	Shift-F5	search-again
Alt-i	overwrite-mode	Ctrl-F5	toggle-case-fold
Alt-j	brief-jump-to-bookmark	Alt-F5	reverse-string-search
Ctrl-k	kill-line	F6	query-replace
Alt-k	kill-to-end-of-line	Shift-F6	replace-again
Alt-l	mark-line-region	Alt-F6	reverse-replace
Alt-m	mark-inclusive-region	F7	record-kbd-macro
Ctrl-n	next-error	Shift-F7	pause-macro
Alt-n	next-buffer	F8	last-kbd-macro
Alt-o	set-file-name	F10	named-command
Alt-p	print-region	Alt-F10	compile-buffer
Ctrl-p	view-process	Ctrl-⟨Enter⟩	brief-open-line
Alt-q	quoted-insert	⟨Esc⟩	abort
Alt-r	insert-file	⟨Del⟩	brief-delete-region
Ctrl-r	argument	⟨End⟩	brief-end-key
Alt-s	string-search	⟨Home⟩	brief-home-key
Ctrl-t	line-to-top	⟨Ins⟩	yank
Alt-t	replace-string	Ctrl-⟨End⟩	end-of-window
Ctrl-u	redo-by-commands	Ctrl-⟨Home⟩	beginning-of-window
Alt-u	undo-by-commands	Ctrl-⟨PgDn⟩	goto-end
Alt-v	show-version	Ctrl-⟨PgUp⟩	goto-beginning
Alt-w	save-file	Alt-⟨Minus⟩	previous-buffer
Ctrl-w	set-want-backup-file	Ctrl-⟨Minus⟩	kill-buffer
Alt-x	exit	Ctrl-⟨Bksp⟩	backward-kill-word
Ctrl-x	write-files-and-exit	Shift-⟨Home⟩	to-left-edge
Alt-z	push	Shift-⟨End⟩	to-right-edge
Ctrl-z	zoom-window	Num +	brief-copy-region
		Num −	brief-cut-region
		Num *	undo-by-commands

Figure 4.9: Epsilon's key map for Brief emulation.

CUA Binding	Epsilon Binding	Command Name
Ctrl-a	Ctrl-x h	mark-whole-buffer
Ctrl-c	Alt-w	copy-region
Ctrl-f	Ctrl-s	incremental-search
Ctrl-h	Alt-r	query-replace
Ctrl-k ...	Ctrl-c ...	(prefix key: see below)
Ctrl-n		new-file
Ctrl-o	Ctrl-x Ctrl-f	find-file
Ctrl-p	Alt-F9	print-buffer
Ctrl-v	Ctrl-y	yank (“paste”)
Ctrl-w ...	Ctrl-x ...	(prefix key: see below)
Ctrl-x	Ctrl-w	kill-region (“cut”)
Ctrl-z	F9	undo
Alt-a	Ctrl-z	scroll-up
Alt-z	Alt-z	scroll-down
Alt-o	Ctrl-x h	mark-paragraph
⟨Escape⟩	Ctrl-g	abort
F3	Ctrl-s Ctrl-s	search-again
⟨Home⟩	Ctrl-a	beginning-of-line
⟨End⟩	Ctrl-e	end-of-line

Figure 4.10: CUA Key Assignments

The `show-variable` command prompts for the name of the variable you want to see, then displays its value in the echo area. The same restrictions on variable types apply here as to `set-variable`. The command includes both user and system variables when it completes on variable names.

The `edit-variables` command in the non-GUI versions of Epsilon lets you browse a list of all variables, showing the current setting of each variable and the help text describing it, as you move through the list. You can use the arrow keys or the normal movement keys to move around the list, or begin typing a variable name to have Epsilon jump to that portion of the list. Press ⟨Enter⟩ to set the value of the currently highlighted variable, then edit the value shown using normal Epsilon commands. To exit from `edit-variables`, press ⟨Esc⟩ or Ctrl-g. With a numeric argument, the command includes system variables in its list.

If Epsilon is set to provide help via a method that permits setting variables directly from the help system (currently just the 32-bit Windows version can do this and only when using WinHelp), the `edit-variables` command behaves differently. It uses the help system to display a list of variables. After selecting a variable, press the Set button to alter its value.

Some Epsilon variables have a different value in each buffer. These *buffer-specific* variables take on a potentially different value each time the current buffer changes. Each buffer-specific variable also has a default value. Whenever you create a new buffer, you also automatically create a new copy of the buffer-specific variable as well. The value of this buffer-specific variable is initially this default value. In Epsilon’s EEL extension language, you can define a buffer-specific variable by using the `buffer` storage class specifier, and give it a default value by initializing it like a regular variable.

Just as Epsilon provides buffer-specific variables, it also provides *window-specific* variables. These have a different value for each window. Whenever you create a new window, you automatically create a new copy of the window-specific variable as well. When you split a window in two, both windows initially have the same values for all their window-specific variables. Each window-specific variable also has a default value.

Epsilon uses the default value of a window-specific variable when it creates its first tiled window while starting up, and when it creates pop-up windows. You define a window-specific variable in EEL with the window storage class specifier, and you may give it a default value by initializing it like a regular variable.

If you ask the `set-variable` command to set a buffer-specific or window-specific variable, it will ask you if you want to set the value for only the current buffer (or window), or the default value, or both. You also can have it set the value in all buffers (or windows).

Variables retain their values until you exit Epsilon, unless you make the change permanent by putting the change in your `einit.ecm` file or by using the `write-state` command, described on page 162. Those methods generally save only the default value for buffer-specific and window-specific variables. They don't save the instantiated values of the variable for each buffer or window, since the buffers and windows themselves aren't listed in an `einit.ecm` or state file. Session files, which do list individual buffers and windows, also record selected buffer-specific and window-specific variables.

The `show-variable` command will generally show you both the default and current values of a buffer-specific or window-specific variable. For string variables, though, the command will ask which you want to see.

The `create-variable` command lets you define a new variable without using the extension language. It asks for the name, the type, and the initial value.

You can delete a variable, command, macro, subroutine, or color scheme with the `delete-name` command, or rename one with the `change-name` command. Neither of these commands will affect any command or subroutine in use at the time you try to alter it.

Summary:	F8	<code>set-variable</code>
	Ctrl-F8	<code>show-variable</code>
		<code>set-any-variable</code>
		<code>edit-variables</code>
		<code>create-variable</code>
		<code>delete-name</code>
		<code>change-name</code>

4.12.5 Saving Changes to Bindings and Variables

Epsilon can save any new bindings you have made and any macros you have defined for future editing sessions. Epsilon uses two kinds of files for this purpose, state files and command files (such as the `einit.ecm` file Epsilon normally uses to save your customizations).

The next section describes command files such as the `einit.ecm` file. For most users, the `einit.ecm` file is the best way to save customizations, but users with a large number of customizations may want to take advantage of the improved startup speed possible by using a state file instead.

Both methods can save bindings, macros, and other sorts of customizations, but they differ in many respects:

- A state file contains commands, macros, variables, and bindings. A command file can contain macros, many types of variables, and bindings, but it can't contain commands written in Epsilon's extension language. (It can contain requests to load other files containing such commands, though.)
- When Epsilon writes a state file, all currently defined commands, macros and variables go into it. A command file contains just what you put there.

- Epsilon can only read a state file during startup. It makes the new invocation of Epsilon have the same commands as the Epsilon that performed the `write-state` command that created that state file. By contrast, Epsilon can load a command file at any time.
- A command file appears in a human-readable format, so you can edit it as a normal file. By contrast, Epsilon stores a state file in a binary format. To modify a state file, you read it into a fresh Epsilon, use appropriate Epsilon commands (like `bind-to-key` to change bindings), then save the state with the `write-state` command.
- Epsilon can read a state file much faster than a command file.
- Binary state files from one release of Epsilon usually aren't compatible with state files from a different major release. Command files are.
- To remove a particular customization from a command file, just delete or comment out its line. Removing a customization from a state file is more complicated, because Epsilon doesn't normally maintain the original setting of a variable, or the original definition of an EEL command you've redefined. You can use `delete-name` to delete a macro, or explicitly set a variable back to its original setting by looking that up in the documentation, but undoing some customizations requires returning to the original state file and reapplying just the customizations you want, via `list-customizations`.

The `write-state` command on `Ctrl-F3` asks for the name of a file, and writes the current state to that file. The file name has its extension changed to `“.sta”` first, to indicate a state file. If you don't provide a name, Epsilon uses the name `“epsilon-v14.sta”`, the same name that it looks for at startup. (The state file name includes Epsilon's major version number.) You can specify another state file for Epsilon to use at startup with the `-s` flag.

By default, when you write a new state file, Epsilon makes a copy of the old one in a file named `ebackup.sta`. You can turn backups off by setting the variable `want-state-file-backups` to 0, or change the backup file name by modifying the `state-file-backup-name` template. See page 122 for information on templates.

Epsilon's default state file sits in its main directory. (In Windows, this is normally under `\Program Files`.) If you write a customized state file, it will go in your customizations directory (see page 13), and Epsilon will read it instead of the default state file. You can use Epsilon's `-s` flag to start Epsilon with its default state file, ignoring your customized one, by running it as `“epsilon -s original”`.

It's a good idea to keep all your customizations in one place, either a state file or an `einit.ecm` file. If you have all your customizations in a state file and want to instead store them all in an `einit.ecm` file, run `list-customizations`, and then delete or rename the customized state file in your customizations directory.

If you have all your customizations in an `einit.ecm` command file and want to instead store them all in a state file, just run Epsilon, letting it load your `einit.ecm`, run `write-state`, and delete or rename your customized `einit.ecm` file.

If you customize Epsilon using an `einit.ecm` file, Epsilon will start up by reading its default state file, which contains its standard settings. Then it will load your customizations from your `einit.ecm` file.

If you customize Epsilon using a state file, Epsilon will read your customized state file instead of the default one.

(If you customize Epsilon using both methods, Epsilon will read your customized state file, then load customizations in your `einit.ecm` file on top. This is confusing, which is why we don't recommend this arrangement.)

The recommended method for saving customizations is to save them all in your `einit.ecm` file, and never write a customized state file.

Command File Syntax

Epsilon's command files appear in a human-readable format, so you can easily modify them. Parentheses surround each command. Inside the parentheses appear a command name, and optionally one or more arguments. The command can be one of several special commands described in the next section, or most any EEL subroutine. See the next section for details.

Each argument can be either a number, a string, or a key list (a special type of string). Spaces separate one argument from the next. Thus, each command looks something like this:

```
(command-name "first-string" "second-string")
```

You can include comments in a command file by putting a semicolon or hash sign ('#') anywhere an opening parenthesis may appear. Such a comment extends to the end of the line. You cannot put a comment inside a string.

For numbers, you can include bases using a prefix of "0x" for hexadecimal, "0o" for octal, or "0b" for binary, or use an EEL-style character constant like 'X' or '\n'. For strings, quote each " or \ character with a \, as in EEL or C.

A few commands such as `define-macro` take a list of one or more keys; these use the syntax of strings, but with some additional rules. Most characters represent themselves, control characters have a "C-" before them, alt characters have an "A-", and function keys have an "F-" followed by the number of the function key. Cursor keys appear in a notation like <Home>. See page 172 for details.

Put a \ before any of these sequences to quote them. For instance, if a macro should contain a Ctrl-f key, write "C-F". If a macro should contain the three characters C hyphen F, write "\C-F". The characters you have to quote are <, ", and \, plus some *letter-* sequences. Thus, the DOS file name \job\letter.txt in a string looks like \\job\\letter.txt. Do not put extra spaces in command file strings that represent keys. For example, the string "C-X F" represents "C-X <Space> F", not "C-X F" with no <Space>.

You can also use the special key syntax <!cmdname> in a keyboard macro to run a command *cmdname* without knowing which key it's bound to. For example, <!find-file> runs the find-file command. When you define a keyboard macro interactively and invoke commands from the menu bar or tool bar, Epsilon will use this syntax to define them, since there may be no key sequence that invokes the specified command.

In addition to the above command syntax with commands inside parentheses, command files may contain lines that define variables, macros, key tables, bindings and color schemes. Epsilon understands all the different types of lines generated by `list-all`, `list-customizations`, `import-customizations`, `list-colors` and similar commands. When Epsilon records customizations in your `einit.ecm` file, it uses this line-by-line syntax for many types of customizations.

Besides listing variables, macros, key tables, and bindings, commands like `list-all` also create lines that report that a particular command or subroutine written in Epsilon's EEL extension language exists. These lines give the name, but not the definition, since these are normally defined in a separate .e file. Typically, an earlier line in the file will have loaded the .e file that defined the command. When Epsilon reads a command line and encounters a line giving the name of a command or EEL subroutine, it simply checks to make sure that a command or subroutine with the given name exists. If not, it reports an error. Epsilon does the same thing with variables that have complicated types (pointers or structures, for example). These also should be defined in .e files.

Command File Examples

One common command in command files is *bind-to-key*. For `bind-to-key`, the first string specifies the name of some Epsilon command, and the second string represents the key whose binding you wish to modify, in a

format we'll describe in detail in a moment. For instance, the following command binds the command `show-matching-delimiter` to `}`:

```
; This example binds show-matching-delimiter to the
; } character so that typing a } shows the matching
; { character.
(bind-to-key "show-matching-delimiter" "{")
```

Unlike the regular command version, `bind-to-key` in a command file can unbind a prefix key. Say you want to make `Ctrl-x` no longer function as a prefix key, but instead have it invoke `down-line`. If, from the keyboard, you typed `F4` to invoke `bind-to-key`, supplied the command name `down-line`, and then typed `Ctrl-x` as the key to rebind, Epsilon would assume you meant to rebind some *subcommand* of `Ctrl-x`, and wait for you to type a `Ctrl-k`, for instance, to bind `down-line` to `Ctrl-x Ctrl-k`. Epsilon doesn't know you have finished typing the key sequence. But in a command file, quotes surround each of the arguments to `bind-to-key`. Because of this, Epsilon can tell exactly where a key sequence ends, and you could rebind `Ctrl-x` as above (discarding the bindings available through `Ctrl-x` in the process) by saying:

```
(bind-to-key "down-line" "C-X")
```

In a command file, *define-macro* allows you to define a keyboard macro. Its first string specifies the name of the new Epsilon command to define, and its second string specifies the sequence of keys you want the command to type. The `define-macro` command does not correspond to any single regular Epsilon command, but functions like a combination of `start-kbd-macro`, `end-kbd-macro`, and `name-kbd-macro`.

In a command file, *set-variable* lets you set a variable with any simple type (numeric or string), just like the usual `set-variable` command. It takes a string with the variable name, a value (either a number or a quoted string), and optionally a numeric code that says how to treat buffer-specific or window-specific variables.

With no numeric code, *set-variable* sets the default value of the variable and its value in all non-system buffers (or windows). A numeric code of 0 sets the value only in the current buffer (or window). The value 1 sets only the default, and 2 sets both. The value 3 sets its value in all non-system buffers or windows, as well as the default value (like omitting the code), while 4 includes system ones too. For example:

```
(set-variable "compile-java-cmd" "oldjavac \"%r\"")
(set-variable "delete-hacking-tabs" 2 1)
```

The *set-variable* command can't enlarge character array variables to accommodate a very long definition. Use the variable-setting syntax produced by commands like `list-all` or `list-customizations` for that; it includes a variable length to set a larger size.

The command file command *create-prefix-command* takes a single string, which specifies a key, and makes that key a prefix character. It works just as the regular command of the same name does.

Using the same syntax, you can also call EEL commands and subroutines. For subroutines, any that take only numeric and string arguments should work. Here are some useful ones:

```
(load-eel-from-path "file.e" 2)
(load-from-path "file.b")
```

The first line makes Epsilon search the `EPSPATH` for an EEL extension language source file named `file.e`, then compile and load it. (You can use an absolute path instead, remembering to double any backslashes.) Its second argument 2 makes the `load_eel_from_path()` subroutine stop if there's an error; the value 1 won't complain if the file wasn't found but complains on other errors, and the value 0 suppresses all errors. (Add the value 4 to make `load-eel-from-path` look for the file in the current directory before

checking the EPSPATH.) Put the EEL file in the same directory as `einit.ecm` (your customization directory) and you can use a relative pathname. (In command names in command files, `-` and `_` are the same.)

The second line makes Epsilon search the EPSPATH for the compiled version of that same file and load it; this is faster, but it requires you to manually recompile that file when updating.

You can even run EEL code directly from a command file, using the `do_execute_eel()` subroutine:

```
(do-execute-eel "
  int i=0;
  while (i++ < 10) {
    has_arg = 1;
    start_process();
  }")
```

Epsilon will compile the code and then execute it. This example starts ten concurrent process buffers.

```
(inline-eel "
command show_color()
{
  char *col = name_color_class(get_character_color(point));
  if (col)
    say("Color class at point is %s.", col);
  else
    say("No color class at point.");
}
")
```

To define commands or variables using EEL code from a command file, you can use the `inline_eel()` subroutine. But it's easier and faster to put longer definitions into a separate `.e` file and load it from the command file with `load-eel-from-path`.

Many commands that prompt for some information like a file name aren't directly suitable for including in a command file; they don't know how to accept an argument supplied by a command file and will always prompt. In many cases an EEL subroutine is available that performs a similar task but is suitable for use in a command file. But a command that takes a numeric prefix argument may be used in a command file: put the prefix argument just before the command name, as in `(100 goto-line)`, which goes to line 100.

Use the *change-name* command to rename an existing command, EEL subroutine, variable, keyboard macro, or color scheme. It takes the old name, then the new one.

Here's an example command file:

```
; This macro makes the window below the
; current one advance to the next page.
(define-macro "scroll-next-window" "C-XnC-VC-Xp")
(bind-to-key "scroll-next-window" "C-A-v")

;This macro asks for a file and puts
;it in another window.
(define-macro "split-and-find" "A-Xsplit-window
A-Xredisplay
A-Xfind-file
")
```

The first two lines contain comments. The third line begins the definition of a macro called `scroll-next-window`. It contains three commands. First `Ctrl-x n` invokes **next-window**, to move to the next window on the screen. The `Ctrl-v` key runs **next-page**, scrolling that window forward, and `Ctrl-x p` then invokes **previous-window**, to return to the original window. The fourth line of this example binds this new macro to the `Ctrl-Alt-v` key, so that from then on, typing a ‘v’ with Control and Alt depressed will scroll the next window forward.

The file defines a second macro named `split-and-find`. It invokes three commands by name. Notice that the macro could have invoked two of the commands by key. Invoking by name makes the macro easier to read and modify later. The `redisplay` command shows the action of `split-window` before the `find-file` command prompts the user for a file name.

Building Command Files

Rather than preparing command files according to the rules presented here, you may wish to have Epsilon write parts of them automatically. You can set the `record-customizations` variable so Epsilon does this automatically whenever you define a macro or bind a key, writing the definition to your `einit.ecm` file.

Or you can have Epsilon list all your current customization settings in a file. The `list-customizations` command inserts a list of all your customizations into the customization buffer, which contains your `einit.ecm` file. This includes bindings, macros, color settings, and other changes. The `list-all` command is similar, but lists every setting, even if it’s not a customization (such as variable settings you’ve never customized).

Epsilon also has commands to create individual `bind-to-key` and `define-macro` lines needed to define a specific current binding or macro.

The `insert-binding` command asks you for a key, and inserts a `bind-to-key` command into the current buffer. When you load the command buffer, Epsilon will restore the binding of that key.

The `insert-macro` command creates an appropriate `define-macro` command for a macro whose name you specify, and inserts the command it builds into the current buffer. This comes in handy for editing an existing keyboard macro. You can use the `edit-customizations` command first to switch to the customization buffer, creating it if necessary.

Summary:

`insert-binding`
`insert-macro`
`list-customizations`

4.13 Advanced Topics

4.13.1 Changing Commands with EEL

Epsilon has many built-in commands, but you may want to add new commands, or modify the way some commands work. We used a language called EEL to write all of Epsilon’s commands. You can find the EEL definitions to all of Epsilon’s commands in files ending in “.e”. EEL stands for Epsilon Extension Language.

Before you can load a group of commands from a “.e” file into Epsilon, you must compile them with the EEL compiler. You do this (outside of Epsilon, or in Epsilon’s concurrent process buffer) by giving the command “`eel filename`” where *filename* specifies the name of the “.e” file you wish to compile (with or without the “.e”). The EEL compiler will read the source file and, if it finds no errors, will produce a “bytecode” file with the same first name but with a “.b” extension. A bytecode file contains command,

subroutine, and variable definitions from the source file translated to a binary form that Epsilon can understand. It's similar to a regular compiler's object file.

Once you've compiled the file, the Epsilon `load-bytes` command gets it into Epsilon. This command prompts for a file name, then loads it into Epsilon. You may omit the extension. But an easier method is to load the EEL source file, then compile and load it in one step using the `compile-buffer` command on Alt-F3. See page 154.

Often a new EEL command won't work the first time. Epsilon incorporates a simple debugger to help you trace through the execution of a command. It provides single-stepping by source line, and you can enter a recursive edit level to locate point, display the values of global variables, or run test functions. The debugger takes the following commands:

- <Space>** Step to the next line. This command will trace a function call only if you have enabled debugging for that function.
- S** If the current line calls a function, step to its first line. Otherwise, step to the current function's next line.
- G** Cancel debugging for the rest of this function call and let the function run. Resume debugging if someone calls the current function again.
- R** Begin a recursive edit of the current buffer. You may execute any command, including `show-variable` or `set-variable`. Ctrl-x Ctrl-z resumes debugging the stopped function. (When debugging a function doing input, you may need to type Ctrl-u Ctrl-x Ctrl-z to resume debugging.)
- T** Toggle whether or not the current function should start the debugger when called the next time. Parentheses appear around the word "Debug" in the debug status line to indicate that you have not enabled debugging for the current function.
- +** Enlarge the debug window.
- Shrink the debug window.
- ?** List all debugger commands.

To start the debugger, use the `set-debug` command. It asks for the name of a command or subroutine, providing completion, and toggles debugging for that function. (A zero numeric argument turns off debugging for that function. A nonzero numeric argument turns it on. Otherwise, it toggles.) The `list-debug` command shows which functions have had debugging set.

Compiling a file with the `-s` EEL compiler flag disables debugging for routines defined in that file. See page 397 for information about the EEL command line options, including the `-s` flag.

The `profile` command shows where a command spends its time. When you invoke the `profile` command, it starts a recursive edit level, and collects timing information. Many times each second, Epsilon notes the source file and source line of the EEL code then executing. When you exit from the recursive edit with Ctrl-x Ctrl-z, Epsilon displays the information to you in a buffer.

Epsilon doesn't collect any profiling information on commands or subroutines that you compile with the `-s` EEL flag.

The `list-undefined` command makes a list of EEL functions that are called from some other EEL function, but have no definition. These are typically the result of misspelled function names.

Summary:

`load-bytes`
`set-debug`
`profile`
`list-undefined`

4.13.2 Updating from an Old Version

When you update to a new release of Epsilon, you'll probably want to incorporate any customizations you've made into the new version. You can save customizations in two ways: in a customization file called `einit.ecm` (see page 164) (a simple text file which this version and future versions of Epsilon can read), or in a state file (a binary file that's specific to a particular major release). Prior to version 12, Epsilon only supported the latter method.

If you're updating from Epsilon 12 or later, and you saved your customizations in an `einit.ecm` file, not a state file, the new version of Epsilon should automatically use your customizations. You might have to modify some if they're affected by changes in the new version of Epsilon; read the release notes to see.

Otherwise, if you're updating from the Windows or Unix version of Epsilon version 8 or later, run the `import-customizations` command in your new version. This will transfer your customizations from your previous version and into your `einit.ecm` file for use in your new version. It works by running your previous version of Epsilon in a special way.

Once you've done that, future versions will automatically use your customizations. If you make more customizations, and choose to save them in a state file, not an `einit.ecm` file, you can use the `list-customizations` command to copy your current customizations to an `einit.ecm` file.

The next section explains how to update from an older version where `import-customizations` isn't supported.

One exception to the above is if you've customized Epsilon by writing Epsilon extensions in Epsilon's extension language, EEL. The `list-customizations` and `import-customizations` commands will help you insert references to your EEL source code files into your `einit.ecm` file, but it's possible that some of the built-in functions or subroutines called by your EEL source code files have changed. In that case, you will have to modify your commands to take this into account. See the section on changes from previous versions in the online manual for information on EEL changes.

Updating from Epsilon 7.0 or Older Versions

This section explains how to transfer customizations from a version of Epsilon so old that the `import-customizations` command isn't supported. The `import-customizations` command is available when updating from version 8 or later.

This section also applies if you're updating from a DOS or OS/2 implementation of Epsilon, regardless of its version, since `import-customizations` is not supported for those platforms.

Moving your customizations (such as variable settings, changed bindings, or keyboard macros) from your old version of Epsilon into the new version requires several steps.

- Start the old version of Epsilon as you do normally.
- Run the `list-all` command.
This will make a list of all the variables, bindings, macros, and functions defined in your old version of Epsilon.
- Save the result in a file. We will assume you wrote it to a file named "after".
- You should no longer need the old version of Epsilon, so you can now install the new version in place of the old one if you wish. Or you can install the new version in a separate directory.
- Locate the "changes" subdirectory within Epsilon's main directory.

For each old version of Epsilon, you'll need several files in the steps below. In the description that follows, we will assume that you want to move from Epsilon 7.0 to this version, and will use files with names like `list70.std`. Substitute the correct file name if you have a different version (for example, `list40.std` to upgrade from Epsilon 4).

- Locate the file in the changes subdirectory from the new version of Epsilon with a name like `list70.std`. It resembles the “after” file, but comes from an unmodified copy of that version of Epsilon. We will call this the “before” file. If you have a very old version for which there is no .std file, see page 172 to make one.
- Start the new version of Epsilon. Run the `list-changes` command. It will ask for the names of the “before” and “after” files, and will then make a list of differences between the files, a “changed” file. When it finishes, you will have a list of the changes you made to the old version of Epsilon, in the format used by the `list-all` command. Edit this to remove changes you don't want in the new version, and save it.
- Run the `load-changes` command, and give it the name of the “changed” file from the previous step. It will load the changes into Epsilon. You can define commands, subroutines, and some variables only from a compiled EEL file, not via `load-changes`. If any of these appear in your changed file, Epsilon will add a comment after that line, stating why it couldn't make the change.
- Use the `list-customizations` command to add your customizations to Epsilon's `init.ecm` file.

Note that this procedure will not spot changes made in .e files, only those made to variables, bindings or macros. It will notice if you have defined a new command, but not if you have modified an existing command.

The above procedure uses several commands. The `list-all` command lists the current state of Epsilon in text form, mentioning all commands and subroutines, and describing all key bindings, macros, and variables. The `list-changes` command accepts the names of the “before” and “after” files produced by `list-all`, and runs the `compare-sorted-windows` command on them to make a list of the lines in “after” that don't match a line in “before”.

Finally, the `load-changes` command reads this list of differences and makes each modification listed. It knows how to create variables, define macros, and make bindings, but it can't transfer extension-language commands. You'll have to use the new EEL compiler to incorporate any EEL extensions you wrote.

Updating from Epsilon 4.0 or Older Versions

If you're updating from a version of Epsilon before 4.0, you'll have to make several files before updating. You will need your old version of Epsilon (including the executable program files for Epsilon and EEL), the state file you've been using with it (typically named `epsilon.sta`), and the original state file that came with that version of Epsilon (which you can find on your old Epsilon distribution disk). You'll also need the file `list-all.e`, included with the new version of Epsilon. First, read the comments in the file `list-all.e` and edit it as necessary to match your version. Then compile it with the old EEL compiler. This will create the bytecode file `listversion.b`. Start your old version of Epsilon with its original state file, using a command like `epsilon -s\oldver\epsilon`, and load the bytecode file you just created, using the `load-bytes` command on the F3 key. Now save the resulting list in a file named “before”. Then start your old version of Epsilon again, this time with your modified state file, and load the bytecode file `listversion.b` again. Now save the resulting list in a file named “after”. Next, start the new version of Epsilon, read in the “before” file, and sort using the `sort-buffer` command, and write it back to the “before” file. You can now continue with the procedure above, running the `list-changes` command and providing the two files you just created.

<Ins>	<Insert>		
<End>			
<Down>			
<PgDn>	<PageDn>	<PgDown>	<PageDown>
<Left>			
<Right>			
<Home>			
<Up>			
<PgUp>	<PageUp>		
	<Delete>		

Figure 4.11: Names Epsilon uses for the cursor keypad keys.

If we didn't provide a .std file for your version of Epsilon, and you're running Epsilon 4.0 or later, here's how to make one. You will need your old version of Epsilon, the state file you've been using with it (typically named epsilon.sta), and the original state file that came with that version of Epsilon (which you can find on your old Epsilon distribution disk). Start your old version of Epsilon with its original state file, using a command like `epsilon -s\oldver\epsilon`, and run the `list-all` command. Now save the resulting list in a file named "before". Then start your old version of Epsilon again (just as you normally do) using the state file that contains the changes you've made, and run the `list-all` command again. Now save the resulting list in a file named "after". Next, start the new version of Epsilon, read in the "before" file, and sort using the `sort-buffer` command, and write it back to the "before" file. You can now continue with the procedure above, running the `list-changes` command and providing the two files you just created.

Summary:

list-all
list-changes
load-changes

4.13.3 Keys and their Representation

This section describes the legal Epsilon keys, and the representation that Epsilon uses when referring to keys and reading command files. The key representation used when writing extension language programs appears on page 557.

Epsilon recognizes hundreds of distinct key combinations you can type on the keyboard (including control and alt keys). You can bind a command to each of these keys. Each key can also function as a prefix key, allowing even more key combinations. By default, Ctrl-x and Ctrl-c serve as prefix keys.

First, the keyboard provides the standard 128 ASCII characters. All the white keys in the central part of the PC keyboard, possibly in combination with the Shift and Control keys, generate ASCII characters. So do the <Esc>, <Backspace>, <Tab>, and <Enter> keys. They generate Control [, Control H, Control I, and Control M, respectively. Depending upon the national-language keyboard driver in use, there may be up to 128 additional keys available by pressing various combinations of Control and AltGr keys, for a total of 256 keys.

You can get an additional 256 keys by holding down the Alt key while typing the above keys. In Epsilon, you can also enter an Alt key by typing an <Esc> before the key. Similarly, the Control-~ key says to interpret the following key as if you had held down the Control key while typing that key.

N-<Ins>	N-<Insert>	N-0	
N-<End>	N-1		
N-<Down>	N-2		
N-<PgDn>	N-<PageDn>	N-<PgDown>	N-<PageDown> N-3
N-<Left>	N-4		
N-5			
N-<Right>	N-6		
N-<Home>	N-7		
N-<Up>	N-8		
N-<PgUp>	N-<PageUp>	N-9	
N-	N-<Delete>	N-.	

Figure 4.12: Numeric keypad key names recognized and displayed by Epsilon.

If you want to enter an actual `<Esc>` or `Control-~` instead, type a `Control-Q` before it. The `Ctrl-q` key “quotes” the following key against special interpretations. See page 158.

In command files and some other contexts, Epsilon represents Control keys by `C-<char>`, with `<char>` replaced by the original key. Thus `Control-t` appears as `C-T`. The case of the `<char>` doesn’t matter for control characters when Epsilon reads a command file, but the `C-` must appear in upper case. The Delete character (ASCII code 127) appears as `C-?`. Note that this has nothing to do with the key marked “Del” on the PC keyboard. The Alt keys appear with `A-` appended to the beginning of their usual symbol, as in `A-f` for `Alt-f` and `A-C-h` for `Alt-Control-H`.

Epsilon represents function keys by `F-1`, `F-2`, . . . `F-63`. The `F` must appear in upper case. You can also specify the Shift, Control, and Alt versions of function keys, in any combination. In a command file, you specify the Shift, Control, and Alt versions with a prefix of `S-`, `C-`, or `A-`, respectively. For example, Epsilon refers to the key you get by holding down the Shift and Alt keys and pressing the `F8` key as `A-S-F-8`.

Keys on the cursor keypad work in a similar way. Epsilon recognizes several synonyms for these keys, as listed in figure 4.11. Epsilon generally uses the first name listed, but will accept any of the names from a command file.

Epsilon normally treats the shifted versions of these keys (and others) as synonyms for the unshifted versions. When you press `Shift-<Left>`, Epsilon runs the command bound to `<Left>`. The commands bound to most of these keys then examine the Shift key and decide whether to begin or stop selecting text. (Holding down the shift key while using the cursor keys is one way to select text in Epsilon.)

Epsilon refers to the numeric keypad keys with the names given in figure 4.12.

In a command file, you can also represent keys by their conventional names, by writing `<Newline>` or `<Escape>`, or by number, writing `<#0>` for the null character `^@`, for example. Epsilon understands the same key names here as in regular expression patterns (see figure 4.3 on page 68).

Macros defined in command files may also use the syntax `<!cmdname>` to run a command `cmdname` without knowing which key it’s bound to. For example, `<!find-file>` runs the `find-file` command. When you define a keyboard macro interactively and invoke commands from the menu bar or tool bar, Epsilon will use this syntax to define them, since there may be no key sequence that invokes the specified command.

Several keys on the PC keyboard act as synonyms for other keys: the grey keys `*`, `—`, and `+` by the numeric keypad, and the `<Backspace>`, `<Enter>`, `<Tab>`, and `<Esc>` keys, for example. The first three act as synonyms for the regular white ASCII keys, and the other four act as synonyms for the Control versions of `‘H’`, `‘M’`, `‘I’` and `‘[’`, respectively. Epsilon normally translates these keys to their synonyms automatically,

Specific Key	Becomes Generic Key	
<code><NumPlus></code>	<code>+</code>	
<code><NumMinus></code>	<code>—</code>	
<code><NumStar></code>	<code>*</code>	
<code><NumSlash></code>	<code>/</code>	
<code><NumEqual></code>	<code>=</code>	
<code><EnterKey></code>	<code><Enter></code>	(on main keyboard)
<code><NumEnter></code>	<code><Enter></code>	(on numeric keypad)
<code><BackspaceKey></code>	<code><Backspace></code>	
<code><TabKey></code>	<code><Tab></code>	
<code><EscapeKey></code>	<code><Esc></code>	
<code><Spacebar></code>	<code><Space></code>	

Figure 4.13: Some keys that are synonyms for others.

and uses the binding of the synonym, but you can also bind them separately if you prefer, using the specific key names shown in figure 4.13.

Mouse Keys

When you use the mouse, Epsilon generates a special key code for each mouse event and handles it the same way as any other key. (For mouse events, Epsilon also sets certain variables that indicate the position of the mouse on the screen, among other things. See page 559.)

<code>M-<Left></code>	<code>M-<LeftUp></code>	<code>M-<DLeft></code>	<code>M-<Move></code>
<code>M-<Center></code>	<code>M-<CenterUp></code>	<code>M-<DCenter></code>	
<code>M-<Right></code>	<code>M-<RightUp></code>	<code>M-<DRight></code>	
<code>M-<Button4></code>	<code>M-<Button4Up></code>	<code>M-<DButton4></code>	
<code>M-<Button5></code>	<code>M-<Button5Up></code>	<code>M-<DButton5></code>	

Epsilon uses the above names for mouse keys when it displays key names in help messages and similar contexts. `M-<Left>` indicates a click of the left button, `M-<LeftUp>` indicates a release, and `M-<DLeft>` a double-click. See page 175 before binding new commands to these keys.

Epsilon doesn't record mouse keys in keyboard macros. Use the equivalent keyboard commands when defining a macro.

There are several "input events" that Epsilon records as special key codes. Their names are listed below. See page 565 for information on the meaning of each key code.

<code>M-<MenuSel></code>	<code>M-<HScroll></code>	<code>M-<WinHelpReq></code>	<code>M-<LoseFocus></code>
<code>M-<Resize></code>	<code>M-<DragDrop></code>	<code>M-<Button></code>	
<code>M-<VScroll></code>	<code>M-<WinExit></code>	<code>M-<GetFocus></code>	

Under Windows, Epsilon displays a tool bar. The `toggle-toolbar` command hides or displays the tool bar. To modify the contents of the tool bar, see the definition of the `standard-toolbar` command in the file menu.e, and the description of the tool bar primitive functions starting on page 525.

The `invoke-windows-menu` command brings up the Windows system menu. `Alt-⟨Space⟩` is bound to this command. If you bind this command to an alphabetic key like `Alt-p`, it will bring up the corresponding menu (the Process menu, in this example).

In a typical Windows program, pressing and releasing the `Alt` key without pressing any other key moves to the menu bar, highlighting its first entry. Set the variable `alt-invokes-menu` to one if you want Epsilon to do this. The variable has no effect on what happens when you press `Alt` and then press another key before releasing `Alt`: this will run whatever command is bound to that key. If you want `Alt-e`, for example, to display the Edit menu, you can bind the command `invoke-windows-menu` to it.

Summary:

`toggle-toolbar`
`invoke-windows-menu`

4.13.4 Customizing the Mouse

You can rebind the mouse buttons in the same way as other keys using the `bind-to-key` command, but if, for example, you rebind the left mouse button to `copy-region`, then that button will copy the region from point to mark, regardless of the location of the mouse. Instead, you might want to use the left button to select a region, and then copy that region. To do this, leave the binding of the left mouse button alone, and instead define a new version of the `mouse-left-hook` function. By default, this is a subroutine that does nothing. You can redefine it as a keyboard macro using the `name-kbd-macro` command. Epsilon runs this hook function after you release the left mouse button, if you've used the mouse to select text or position point (but not if, for example, you've clicked on the scroll bar).

Normally Epsilon runs the `mouse-select` command when you click or double-click the left mouse button, and the `context-menu` command when you click or double-click the right mouse button. Epsilon runs the `mouse-move` command when you move the mouse; this is how it changes the mouse cursor shape or pops up a scroll bar or menu bar when the mouse moves to an appropriate part of the screen, in some environments.

It's possible to bind a mouse button to the `mouse-to-tag` function, which gets the name of an identifier where you clicked and then goes to its definition via tags (see page 50). Or you can use the right-click context menu's Go to Definition command for this. Both `mouse-select` and `mouse-to-tag` run the appropriate hook function for the mouse button that invoked them, whenever you use the mouse to select text or position point. The hook functions for the other two mouse buttons are named `mouse-right-hook` and `mouse-center-hook`. You can redefine these hooks to make the mouse buttons do additional things after you select text, without having to write new commands using the extension language. (The `context-menu` command doesn't call any hook function.)

By default, the center mouse button runs the command `mouse-center`, which in turn calls the `mouse-pan` command to make the mouse scroll or pan. Setting the `mouse-center-yanks` variable makes it perform a different action. Some settings make it call the `mouse-yank` command, to have the middle mouse button yank text from the clipboard (a traditional function under Unix).

Mice with up to five buttons are supported on Windows and X11. By default, buttons 4 and 5 page up and down.

In Windows, clicking in Epsilon's window when another program's window has the focus sets point to the place you clicked. If you would like clicking to switch to Epsilon but not change point, you can define the Epsilon configuration variable (see page 10) `NOFOCUSCLICK` in the registry.

Summary:

`M-⟨Left⟩` `mouse-select`
`M-⟨Right⟩` `context-menu`

M-⟨Center⟩	mouse-center
M-⟨Move⟩	mouse-move
M-⟨Button4⟩	next-page
M-⟨Button5⟩	previous-page
	mouse-pan
	mouse-yank
	mouse-to-tag

4.14 Miscellaneous

You can use the `eval` command to quickly evaluate an arbitrary EEL expression, or do simple integer-only math. By default, the command displays the result; use a numeric prefix argument and it will insert the result in the current buffer. You can append a `;` character followed by a `printf`-style format specification, and Epsilon will use that format for the result. For example, `403*2;x` displays the value 806 converted to hexadecimal, `0x03B5;k` displays the name of Unicode character U+03B5, and `"simple";.3s` displays `"sim"`.

Similarly, the `execute-eel` command executes a line of EEL code that you type in.

The command `narrow-to-region` temporarily restricts your access to the current buffer to the region between the current values of `point` and `mark`. Epsilon hides the portion of the buffer outside this region. Searches will only operate in the narrowed region. While running with the buffer narrowed, Epsilon considers the buffer to start at the beginning of the region, and end at the end of the region. However, if you use a file-saving command with the buffer narrowed in this manner, Epsilon will write the entire file to disk. To restore normal access to the buffer, use the `widen-buffer` command.

Under Windows, you can set Epsilon's key repeat rate with the `key-repeat-rate` variable. It contains the number of repeats to perform in each second. Setting this variable to 0 lets the operating system or keyboard determine the repeat rate, as it does outside of Epsilon. Epsilon never lets repeated keys pile up; it ignores automatically repeated keys when necessary.

Summary:

narrow-to-region
widen-buffer
eval
execute-eel

Chapter 5

Alphabetical Command List



abort	Ctrl-g	Abort the currently executing command.
-------	--------	--

This special command causes a currently executing command to stop, if possible. It cancels any executing macros, and discards any characters you may have typed that Epsilon hasn't read. Use the `set-abort-key` command to change the abort key.

about-epsilon	Show Epsilon's version number and operating system.
---------------	---

align-by-tab Interpret Tab characters to make columns line up.

This command takes text containing tab characters (such as text copied from a spreadsheet) and reformats it using spaces and tabs so the text appears visually in columns, making each column just wide enough for its contents.

align-region	Ctrl-c Ctrl-a	Change spacing to line up similar things.
--------------	---------------	---

This command looks for strings that occur on more than one line in the region, then changes the spacing before them to make them line up. By default, it does this for “=” characters and comments. Modes can define additional strings to align. C mode also aligns the names of variables in simple variable definitions and the definitions of `#define` lines. The `align-region-rules` variable controls what will be aligned.

With a numeric argument, this command does manual alignment. It prompts for a regular expression pattern. The pattern must consist of two subpatterns, each enclosed in parentheses, such that #1 and #2 would match the subpatterns in a regular expression replacement operation. For each line containing a match for the pattern, the spacing at the boundary between the two subpatterns will be modified. The new spacing will be based on the widest match for #1.

To align the semicolons on each line, for example, you could use the pattern `() (;)`. Epsilon uses the first match of the pattern on each line, skipping lines that don't contain a match.

alt-prefix	ESC	Interpret the next key as an Alt key.
------------	-----	---------------------------------------

This command reads a character from the keyboard, then runs the command bound to the Alt version of that key.

ansi-to-oem Convert buffer's Windows character set to DOS.

Windows programs typically use a different character set than do DOS programs. The DOS character set is known as the DOS/OEM character set, and includes various line drawing characters and miscellaneous characters not in the Windows/ANSI set. The Windows/ANSI character set includes many accented characters not in the DOS/OEM character set. Epsilon for Windows uses the Windows/ANSI character set (with most fonts).

The ansi-to-oem command converts the current buffer from the Windows/ANSI character set to the DOS/OEM character set. If any character in the buffer doesn't have a unique translation, the command warns first, and moves to the first character without a unique translation.

This command ignores any narrowing established by the `narrow-to-region` command. It's only available in Epsilon for Windows.

`back-to-tab-stop` `Shift-⟨Tab⟩` Move back to the previous tab stop.

This command moves point to the left until it reaches a tab stop, a column that is a multiple of the tab size.

If a region is highlighted, Epsilon unindents the region by one tab stop. With a numeric prefix argument, Epsilon unindents by that amount.

This command uses the variable `soft-tab-size` if it's nonzero. Otherwise it uses `tab-size`.

`backward-character` `Ctrl-b` Move point back.

Point moves back one character. Nothing happens if you run this command with point at the beginning of the buffer. If you set `virtual-space` to two, the command never moves to a different line but only left on the current line.

`backward-delete-character` `Ctrl-h` Delete the character before point.

This command deletes the character before point. When given a numeric argument, the command deletes that many characters, and saves them in a kill buffer.

This command can be set to convert tab characters into spaces before deleting, or to delete multiple spaces at once. See the `delete-hacking-tabs` variable.

`backward-delete-word` Brief: `Ctrl-⟨Backspace⟩` Delete the word before point.

The command moves point as in `backward-word`, deleting the characters it passes over. See `backward-kill-word` for a similar command that cuts the text to a kill buffer.

`backward-ifdef` C mode: `Alt-[`, `Alt-⟨Up⟩` Find matching preprocessor line.

This command moves to the previous `#if/#else/#endif` (or similar) preprocessor line. When starting from such a line, Epsilon finds the previous matching one, skipping over inner nested preprocessor lines.

`backward-kill-level` `Alt-⟨Del⟩` Kill a bracketed expression backwards.

The command moves point as in `backward-level`, killing the characters it passes over.

`backward-kill-word` `Ctrl-Alt-h` Kill the word before point.

The command moves point as in `backward-word`, killing the characters it passes over.

`backward-level` `Ctrl-Alt-b` Move point before a bracketed expression.

Point moves backward searching for one of `)`, `}`, or `]`. Then point moves back past the nested expression and positions point before the corresponding left delimiter. (Actually, each mode defines its own list of delimiters. One mode might only recognize `<` and `>` as delimiters, for instance.)

`backward-paragraph` `Alt-[` Go back one paragraph.

Point travels backward through the buffer until positioned at the beginning of a paragraph. Lines that start with whitespace (including blank lines) always separate paragraphs. For information on changing Epsilon's notion of a paragraph, see the `forward-paragraph` command.

backward-sentence Alt-a Go back one sentence.

Point travels backwards through the buffer until positioned at the beginning of a sentence. A sentence ends with a period, exclamation point, or question mark, followed by two spaces or a newline, with any number of closing characters " , ' ,) ,] , between. A sentence also ends at the end of a paragraph. See forward-paragraph.

backward-word Alt-b Go back one word.

Point travels backward until positioned before the first character in some word.

batch-mode Set up for editing Windows Batch files.

This command puts the current buffer in Batch mode, suitable for .bat, .cmd, and .btm files from Windows, DOS, or OS/2.

beginning-of-line Ctrl-a Go to the start of the line.

This command positions point at the beginning of the current line, before the first character.

beginning-of-window Alt-, Go to the upper left corner.

Position point before the first character in the window.

bind-last-macro Ctrl-x Alt-n Put a new macro on a key.

After you define a keyboard macro, you can use this command to assign a name to it and bind it to a key. The command makes up a name for the macro, which you can change, then asks for the key sequence that should invoke it. It combines the functions of the name-kbd-macro and bind-to-key commands.

bind-to-key F4 Put a named command on a key.

This command prompts you for the name of a command, then for a key. Thereafter, pressing that key runs that command.

Certain keys like <Backspace> have both a generic interpretation (Ctrl-h) and a specific one. If you press one of these keys, Epsilon will ask which binding you want to make. Select the generic one unless you want each way of typing the generic key (such as Ctrl plus H, versus the <Backspace> key) to perform a different function.

brief-copy-region Brief: Grey + Copy a highlighted region, saving it.

This command saves a copy of the highlighted region to a kill buffer so you can insert it somewhere else. If no region is highlighted, the command copies the current line.

brief-cut-region Brief: Grey - Delete a highlighted region, saving it.

This command kills the highlighted region, saving a copy of the region to a kill buffer so you can insert it somewhere else. If no region is highlighted, the command kills the current line.

brief-delete-region Brief: Delete a highlighted region without saving.

This command deletes the highlighted region without saving it in a kill buffer. If no region is highlighted, the command deletes the next character in the buffer.

brief-delete-window Brief: F4 Remove one of a window's borders.

This command prompts you to indicate which of the current window's borders you wish to delete. Press an arrow key and Epsilon will delete other windows as needed to remove that window border.

brief-drop-bookmark Brief: Alt-0 ... Alt-9 Remember this location.

This command remembers the current buffer and position, so that you can easily return to it later with **brief-jump-to-bookmark**. Normally, the command looks at the key you pressed to invoke it, to determine which of the ten Brief bookmarks to set. For example, if you press Alt-3 to invoke it, it sets bookmark 3. If you press Alt-0 to invoke it, it sets bookmark 10. When you invoke the command by pressing some other key, it prompts for the bookmark to set.

Brief bookmarks 1–10 are really synonyms for Epsilon bookmarks A–M. You can use Epsilon commands like **list-bookmarks** to see all the bookmarks and select one.

brief-end-key Brief: <End> Go to the end of the line/window/buffer.

This command goes to the end of the current line. When you press it twice in succession, it goes to the end of the current window. When you press it three times in succession, it goes to the end of the current buffer.

brief-home-key Brief: <Home> Go to the start of the line/window/buffer.

This command goes to the start of the current line. When you press it twice in succession, it goes to the start of the current window. When you press it three times in succession, it goes to the start of the current buffer.

brief-jump-to-bookmark Brief: Alt-j Jump to a bookmark.

This command returns to a bookmark previously set with **brief-drop-bookmark**. It prompts for the number of the bookmark you wish to return to.

Brief bookmarks 1–10 are really synonyms for Epsilon bookmarks A–M. You can use Epsilon commands like **list-bookmarks** to see all the bookmarks and select one.

brief-keyboard Load the Brief-style keyboard layout.

This command redefines the keyboard to resemble the key arrangement used by the Brief editor. Use the command **epsilon-keyboard** to return to Epsilon's default keyboard arrangement.

brief-open-line Brief: Ctrl-<Enter> Make a new line below this one.

This command adds a new line after the current one and moves to it.

brief-resize-window Brief: F2 Move the window's border.

This command prompts you to indicate which of the current window's borders you would like to move. Press an arrow key to select one. Then press arrow keys to move the window's border around. Press `<Enter>` when you are satisfied with the window's size. Epsilon will resize other windows as necessary.

brief-split-window Brief: F3 Put a new border inside this window.

This command prompts you to indicate where you would like to create a new window border. Press an arrow key and Epsilon will split off a new window from the current one, with the border between the two in the indicated direction.

browse-current-symbol Browse source code for this symbol without prompting.

Epsilon provides an interface to source code browsing data generated by Microsoft compilers, using a buffer in Browse mode. Use this command to display the data available for a particular symbol, such as a function or variable name.

This is a variation on the `browse-symbol` command. It doesn't prompt for a symbol name, but uses the name at point instead (or the name in the highlighted region, if any). See the description of the `browse-symbol` command for details.

browse-mode Source code browsing interface.

Epsilon provides an interface to source code browsing data generated by Microsoft compilers, using a buffer in Browse mode. See the `browse-symbol` command to create a Browse mode buffer.

In Browse mode, pressing `<Enter>`, `<Space>`, or "e", or double-clicking, displays the item mentioned on the current line; the exact behavior depends on the line's contents. If the line refers to a specific source file name and line number, Epsilon goes to that file and line, in the current window if there's only one window, or in another window if there's more than one.

If the current line refers to a different symbol, in the uses or used-by section of the listing, Epsilon displays the Browse data for that symbol, replacing the current symbol. Press `L` to return to the previous symbol, moving back in symbol history. Press `Ctrl-u L` to display the symbol history and select a previous symbol to display.

For lines in the heading section of the Browse display, Browse mode prompts for a new setting for the item displayed on that line.

browse-set-filter Browse mode: f Filter out some instances in source code browsing.

A Browse mode buffer created by the `browse-symbol` command displays uses of the symbol as a function, variable, macro, and so forth. You can exclude some of these types of uses by setting a filter using this command.

The command shows the current uses permitted. Press one of the indicated keys to toggle a particular type of use. Press `<Enter>` when done and the current listing will be updated.

browse-set-usedby-filter Browse mode: b Filter out some uses in source code browsing.

A Browse mode buffer created by the `browse-symbol` command displays those function, variable, macro, and so forth that use a particular symbol, and those used by the current symbol. You can exclude some of these types of uses by setting a filter using this command.

The command shows the current uses permitted. Press one of the indicated keys to toggle a particular type of use. Press `<Enter>` when done and the current listing will be updated.

browse-symbol Ctrl-`<NumSlash>` Browse source code for this symbol.

Epsilon provides an interface to source code browsing data generated by Microsoft compilers, using a buffer in Browse mode. Use this command to display the data available for a particular symbol, such as a function or variable name.

First you must set up Epsilon for source code browsing by turning on browsing data in your compiler and running the `configure-epsilon` command to install a .DLL file. See page 52 for details.

If you haven't already either run the `select-browse-file` command to select your .bsc browser database file, or used the `select-tag-file` to select a .bsc file, this command will prompt for your .bsc file the first time you run it.

The command then prompts for the name of a symbol and displays its data. Symbol name completion is available. You can put a * at the end of the symbol name to display all symbols starting with the given prefix, or write `classname::` to display all that class's members.

The `#symbols#` buffer it creates is in browse-mode; see that topic for details on its commands. Also see the `browse-current-symbol` command.

bufed Ctrl-x Ctrl-b Manipulate a list of buffers.

This command makes a list of buffers, displays it in a window, and lets you edit it. Alphabetic keys run special bufed commands. The N and P commands go to the next and previous buffers in the list, respectively. The D command deletes the buffer on the current line immediately. It warns you if the buffer has unsaved changes. The `<Space>` or E key selects the buffer on the current line, and the S key writes the buffer named on the current line to its file. Typing 1 makes the window occupy the whole screen, then selects the buffer like E. Typing 2 or 5 splits the window horizontally or vertically, then selects the indicated buffer. Shift-P prints the buffer on the current line.

In a bufed listing, the A, B, F, and I keys make bufed sort the buffer list by last access time, buffer name, file name, or size, respectively. Use the shifted versions of these keys to sort in reverse. Pressing U requests an unsorted buffer list: the newest buffers appear first in the list. M toggles whether modified buffers appear first in the list.

This command does not normally list special buffers such as the kill buffers whose names begin with the "-" character. To list even these buffers, give the bufed command (or a sorting command) a numeric argument.

buffer-grep Search in buffers for a pattern.

This command searches through buffers for a pattern. It prompts for a search string and a buffer pattern; the default is to search only the current buffer. Then it searches in all buffers matching the buffer pattern, listing matching lines in the grep buffer. The grep buffer then appears in the current window.

By default, `buffer-grep` interprets the search string as a regular expression. Press `Ctrl-t` at the search string prompt to toggle regular expression mode. You can also type `Ctrl-w` or `Ctrl-c` to toggle word-mode or case-folding searches, respectively.

The buffer name pattern may contain the wildcard characters `?` to match any single character, `*` to match zero or more characters, a character class like `[^a-zA-Z]` to match any non-alphabetic character, or `|` to separate alternatives.

In `grep` mode, alphabetic keys run special `grep` commands. See the description of the `grep-mode` command for details. Typing `H` or `‘?’` in `grep` mode gives help on `grep` subcommands.

A bit in the `search-in-region` variable makes grepping a single buffer restrict itself to a highlighted region. By default, `Epsilon` searches entire buffers.

`buffer-spell-mode` Toggle per-buffer highlighting of misspelled words.

This command toggles the `Spell` minor mode for the current buffer only. With a nonzero numeric argument, it turns `Spell` mode on; with a zero argument it turns `Spell` mode off. In `Spell` mode, `Epsilon` highlights misspelled words. See the `spell-mode` command to toggle `Spell` minor mode for all buffers with the same major mode as the current buffer (for instance, all `HTML` buffers).

Use the `spell-correct` command to ignore certain words or show suggestions for a misspelled word.

`c-close` C mode: `},)` Self-insert, then fix this line's indentation for C.

`c-colon` C mode: `:` Self-insert, then fix this line's indentation for C.

`c-hash-mark` C mode: `#` Self-insert, then fix this line's indentation for C.

`c-mode` Do automatic indentation for C-like languages.

This command puts the current buffer in C mode, appropriate for editing programs written in any language with a syntax similar to C (such as `EEL`). In C mode, `⟨Enter⟩` indents each new line by scanning previous lines to determine the proper indentation. `⟨Tab⟩` reindents the current line when you invoke it with point inside a line's indentation. With point outside a line's indentation, or when repeated, this command adds more indentation.

By default, the `find-file` command automatically turns on C mode for files that end with `.c`, `.cpp`, `.hpp`, `.cxx`, `.hxx`, `.y`, `.h`, `.java`, `.idl`, `.idl`, `.acf`, `.cs`, `.i`, `.ii`, `.m`, `.mi`, `.mm`, `.mmi`, or `.e`.

`c-open` C mode: `{` Self-insert, then fix this line's indentation for C.

`capitalize-word` `Alt-c` Upper case beginning character.

Point travels forward through the buffer as with `forward-word`. Each time it encounters a run of alpha characters, it converts the first character to upper case, and the remainder to lower case.

For example, if you execute this command with point positioned just before `“wORd”`, it becomes `“Word”`. Similarly, `“wORd_wORd”` becomes `“Word_Word”`.

If the current buffer contains a highlighted region, Epsilon instead capitalizes all the words in the region, leaving point unchanged.

cd **F7** **Change the current directory.**

The `cd` command prompts for the name of a directory, then sets Epsilon's current directory. If you press Alt-e when prompted for the directory name, Epsilon will type in the name of the directory portion of the current file name.

When Epsilon displays a file name (for example, in a buffer's mode line), it usually describes the file relative to this current directory.

Epsilon uses its notion of the current directory when it prompts for a file name and the current buffer has no specific directory associated with it. (This typically happens when the buffer has no associated file name.)

Also, if you remove any pre-typed directory name and type a relative pathname to such a command, Epsilon will interpret what you type relative to the directory set by the `cd` command.

See the `prompt-with-buffer-directory` variable for more information.

center-line **Alt-s** **Center line horizontally.**

This command centers the current line between the first column and the right margin, by changing the line's indentation if necessary. You can set the right margin with the `set-fill-column` command.

center-window **Ctrl-l** **Vertically center the current window.**

This command makes the line containing point appear in the center of the window. With a numeric argument *n*, it makes the line appear on line *n* of the window. Line 0 refers to the top line.

change-code-coloring **Toggle code coloring on or off in this buffer.**

This command toggles between coloring and not coloring the program text in the current buffer by setting the `want-code-coloring` variable. The command has no effect in buffers Epsilon doesn't know how to color.

change-file-read-only **Change the read-only status of a file.**

This command prompts for a file name (default: the current file) and toggles its read-only attribute. Under Unix, it either makes the file unwritable to all, or writable to all (to the extent permitted by the current `umask`). Use Alt-o ! `chmod` for finer control.

change-font-size **Set the font's width and height.**

Under Windows, this command supplements the `set-font` command by providing additional font choices. Some Windows fonts include a variety of character cell widths for a given character cell height. (For example, many of the font selections available in windowed DOS sessions use multiple widths.) Commands like `set-font` utilize the standard Windows font dialog, which doesn't provide any way to select these alternate widths. This command lets you choose these fonts.

`clear-tags` Forget all the tags in the current tag file.

See also the commands `select-tag-file` and `tag-files`.

`comment-region` Ctrl-c Ctrl-r Mark the region as a comment.

This command comments out each line in the highlighted region, using the current mode's comment syntax. With a numeric prefix argument, it uncomments lines in the region instead. If no region is highlighted, it operates on the current line.

`compare-files` Compare files from two directories.

This command prompts for a pair of directories, then shows a list of files that are the same in the two, different, or missing from one or the other directory hierarchy. You can specify a file pattern instead of one or both directories to list only the files that match. With a prefix argument, the command asks for the comparison method and destination buffer.

`compare-files-mode` Examine a list of differences in directories.

In a list of files produced by the `compare-files` command, you can use these keys:

- n** moves to the next file in the list.
- p** moves to the previous file.
- a** visits the current line's file matching the first directory or pattern specified.
- b** visits the current line's file matching the second directory or pattern specified.
- c** opens the current line's file from both the first and second directory or pattern specified in two windows one above the other.
- s** opens the current line's file from both the first and second directory or pattern specified in two windows side by side.
- v or** `<Enter>` displays a visual-diff of the current line's two files.
- d** displays a diff of the current line's two files.
- l** (lowercase L) displays a diff of all the different files in the list, concatenated.
- Shift-A** copies the current line's file from the first directory or pattern specified to the second, prompting before overwriting it.
- Shift-B** copies the current line's file from the second directory or pattern specified to the first, prompting before overwriting it.
- h or ?** gives this help.

`compare-sorted-windows` Find lines missing from the current or next windows.

This command copies all lines that appear in both the current window's buffer, and the next window's buffer, into a buffer named "inboth". It copies other lines to buffers named "only1" and "only2". It assumes that you already sorted the original buffers.

`compare-to-prior-version` Show changes from editing.

This command prompts for an edit count *n*, then compares the current version of the buffer with the version prior to the most recent *n* editing operations, displaying the differences in a new buffer. An "editing operation" means what one use of the `undo-changes` command would undo. If the edit count you enter is zero, the command instead compares the current version of the buffer with the version currently on disk. The comparison uses the format of the `visual-diff` command by default; set the `compare-to-prior-version-style` variable to use a different format.

compare-windows	Ctrl-F2	Find the next difference between the current and next windows.
-----------------	---------	--

This command moves forward from point in the buffers displayed in the current window and the next window. It compares the text in the buffers, stopping when it finds a difference or reaches the end of a buffer, then reports the result.

If repeated, it alternates between finding the next difference and finding the next match (by resynchronizing the buffers).

compile-buffer	Alt-F3	Compile the current buffer as appropriate.
----------------	--------	--

This command tries to compile the current buffer. It uses the compiling command appropriate for the current buffer. For .c files, this is contained in the `compile-c-cmd` variable. For .cpp or .cxx files, this is contained in the `compile-cpp-cmd` variable. For .e files, this is contained in the `compile-eel-cmd` variable. When you compile an EEL file successfully, Epsilon automatically loads the resulting bytecode file.

If the current buffer has no compilation command associated with it, Epsilon will prompt for the appropriate command and record it in the buffer-specific variable `compile-buffer-cmd`. For C, C++, and EEL files, Epsilon automatically sets this to refer to the variables listed above.

Before and after running the compilation command, Epsilon does any mode-specific operations needed, by calling the buffer-specific function pointer variables `pre_compile_hook` and `post_compile_hook`, respectively. An EEL programmer can use these hooks to make Epsilon perform additional actions each time you compile buffers.

The function pointed to by `post_compile_hook` receives a status code returned by the `do_compile()` subroutine, and the number of the buffer to be compiled. See that function's definition in `proc.e` for details. The function pointed to by `pre_compile_hook` receives no parameters. If either variable holds a null pointer, Epsilon doesn't call it.

conf-mode		Set up for editing configuration files.
-----------	--	---

This command sets up generic syntax highlighting suitable for miscellaneous Unix configuration files.

configure-epsilon		Under Windows, set up file associations and similar.
-------------------	--	--

In Epsilon for Windows, this command launches a setup utility that lets you configure file associations, recreate Epsilon's Start Menu entries and other shortcuts, install source code browser support, and modify registration information.

context-help	Shift-F1	Display help on the keyword at point.
--------------	----------	---------------------------------------

This command displays help on the keyword at point. The type of help depends on the current mode. Context-sensitive help is configurable by setting mode-based help rule variables such as `context-help-rule-perl`. See page 102 for details on these variables.

context-menu	Shift-F10	Display a context menu.
--------------	-----------	-------------------------

This command displays a context menu. The right mouse button runs this command.

`copy-to-scratch` `Ctrl-x x` Copy the region to a permanent buffer.

This command copies the text in the region between point and mark. It asks for a letter (or number), then associates that character with the text. Subsequently, you can insert the text by invoking the `insert-scratch` command. See also the commands `kill-region` and `copy-region`.

`count-lines` `Ctrl-x L` Show the number of lines in the buffer.

A message showing the number of lines in the buffer appears in the echo area. The message also gives the line number of the current line, and the length of the file when written to disk. If there is a highlighted region, its line count is displayed as well.

See the `mode-format` variable to display the current line number continuously.

`count-words` `Ctrl-c Ctrl-w` Show the number of words in the buffer.

This command displays the number of words in the current buffer. If there's a highlighted region, it counts only words within the region. With a numeric argument, it prompts for a regular expression pattern and counts instances of that, instead of words.

`create-file-associations` Make Windows run Epsilon to launch certain file types.

You can set up Windows file associations for Epsilon using the `create-file-associations` command. It lets you modify a list of common extensions, then sets up Windows to invoke Epsilon to edit files with those extensions. The files will be sent to an existing copy of Epsilon, if one is running. If not, a new instance of Epsilon will be started.

`create-prefix-command` Define a new prefix key.

This command asks for a key and then turns that key into a prefix key, like `Ctrl-x`.

`create-variable` Define a new EEL variable.

This command lets you define a new variable without using the extension language. It prompts for the name, the type, and the initial value.

`css-mode` Set up for editing CSS files.

This command puts the current buffer in CSS mode, suitable for editing cascading style sheet files used for web pages.

`ctrl-prefix` `Ctrl-^` Interpret the next key as a Control key.

This command reads a character from the keyboard, then executes the command bound to the Control version of that key.

`cua-keyboard` Load the CUA-style keyboard layout.

This command redefines the keyboard to resemble the key arrangement used by typical MS-Windows programs. Use the command `epsilon-keyboard` to return to Epsilon's default keyboard arrangement.

`delete-blank-lines` `Ctrl-x Ctrl-o` Remove blank lines around point.

This command deletes empty lines adjacent to point, or lines that contain only spaces and tabs, turning two or more such blank lines into a single blank line. The command deletes a lone blank line. If you prefix a numeric argument of n , exactly n blank lines appear regardless of the number of blank lines present originally.

With a highlighted region, the command does this at every sequence of one or more blank lines throughout the region.

`delete-character` `Ctrl-d` Delete the character after point.

If you prefix a numeric argument, the command deletes that many characters, and saves them in a kill buffer. If invoked immediately after a kill command, `delete-character` will store the deleted character(s) in the same kill buffer that the kill command used (but see the `delete-options` variable).

`delete-current-line` Brief: `Alt-d` Delete the current line.

This command deletes the entire current line, including any newline at its end.

`delete-horizontal-space` `Alt-\` Delete whitespace near point.

This command deletes spaces and tabs surrounding point.

`delete-matching-lines` Delete lines containing a regex pattern.

This command prompts for a regular expression pattern. It then deletes all lines below point in the current buffer that contain the pattern. While you type the pattern, `Ctrl-w` enables or disables word searching, restricting matches to complete words. `Ctrl-t` enables or disables regular expression searching, in which the search string specifies a pattern (see `regex-search` for rules). `Ctrl-c` enables or disables case-folding.

`delete-name` Delete a function, variable, etc.

This command prompts you for the name of a command, subroutine or variable, with completion, and then tries to delete the item.

`delete-rectangle` Delete the characters in the current rectangle.

This command removes from the current buffer the characters in the rectangular area between point and mark. Unlike the `kill-rectangle` command, this command does not copy the characters to a kill buffer. Also see the `append-next-kill` command.

`delete-region` Delete the region without saving it.

This command removes the characters between point and mark from the buffer without putting them in a kill buffer.

`delete-to-end-of-line` Delete the remaining characters on this line.

describe-command F1 C Give help on the named command.

This command prompts you for a command name, then displays a description of that command along with its current bindings (if any).

describe-key F1 K Give help on the key.

This command prompts you for a key, then displays a description of the command bound to that key (if any).

describe-variable F1 R Display help on a variable.

This command prompts for the name of variable. Then it displays the documentation for that variable.

dialog-regex-replace Replace using a dialog.

This command displays the Replace dialog, which you can use to find and replace text in the buffer. The dialog is initialized so that the Regular Expression box is checked.

dialog-replace Replace using a dialog.

This command displays the Replace dialog, which you can use to find and replace text in the buffer.

To match special characters, turn on regular expression searching, and then enter them using the syntax <Tab> or <#13>. In replacement text, use the similar syntax #<NewLine>.

dialog-reverse-search Search backwards using dialog.

This command displays a Find dialog initialized to search backwards.

dialog-search Search using the Find dialog.

This command displays a Find dialog, which you can use to search for text in the buffer.

To match special characters, turn on regular expression searching, and then enter them using the syntax <Tab> or <#13>.

diff List differences between current and next windows.

Make a list of all differences between the buffers in the current and next windows in a buffer named #diff#. The list shows what lines you would have to remove from or add to the first buffer to make it identical to the second buffer. With a numeric argument, the command prompts for the name of a buffer to hold the results.

dired Ctrl-x d Edit the contents of a directory.

The command dired (for directory edit) allows you to conveniently peruse the contents of a directory, examining the contents of files and, if you wish, selecting some for deletion, copying, or moving.

The command prompts for the name of a directory or a file pattern. By default, it uses the current directory. It then displays a buffer in the current window, with contents similar to what the operating system command “dir” would display. Each line of the dired buffer contains the name of a file and information about it.

In dired mode, alphabetic keys run special dired commands. See the description of the dired-mode command for details. Typing H or ‘?’ in dired mode gives help on dired subcommands.

dired-mode

Edit a directory of file names.

A dired (directory edit) buffer lists the contents of a directory. In a dired buffer, you can use these keys:

- n** moves to the next entry in the list.
- p** moves to the previous entry.
- d** flags a file (or empty directory) that you wish to delete by placing a ‘D’ before its name.
- c** marks a file for copying.
- m** marks a file for moving (renaming).
- u** removes any flags from the file listed on the current line.
- x** actually deletes, copies, or moves the files. Epsilon will ask for the destination directory into which the files are to be copied or moved, if any files are so marked. If there is only one file to copy or move, you can also specify a file name destination, so you can use the command for renaming files. Epsilon prompts for a single destination for all files to be copied, and another for all files to be moved. If any files are marked for deletion, Epsilon will ask you to confirm that you want to delete the files.
- e or** `<Space>` **or** `<Enter>` lets you examine the contents of a file. It invokes the find-file command on the file, making the current window display this file instead of the dired buffer. After examining a file, you can use the select-buffer command (`Ctrl-x b`) to return to the dired buffer. Press `<Enter>` when prompted for the buffer name and the previous buffer shown in the current window will reappear (in this case, the dired buffer). Applied to a directory, the E command does a dired of that directory.
- Shift-E** examines the contents of a file or directory like `<Enter>`, deleting the current dired buffer first.
- lowercase L** creates a live link. First Epsilon creates a second window, if there’s only one window to start with. (Provide a numeric argument to get vertical, not horizontal, window splitting.) Then Epsilon displays the file named on the current dired line in that window, in a special live link buffer. As you move around in the dired buffer, the live link buffer will automatically update to display the current file. Delete the live link buffer or window, or show a different buffer there, to stop the live linking.
- v** runs the “viewer” for that file; the program assigned to it according to Windows file associations. For executable files, it runs the program. For document files, it typically runs the Windows program assigned to that file extension. In Epsilon for Unix, it tries to display the file using the KDE, Gnome, or macOS view setting for that type of file, by calling an `epsilon-viewer` script you can customize.
- t** displays the MS-Windows properties dialog for that file or directory. For a directory, this lets you view the size of its contents.
- @** toggles the read-only property of the current file. Under Mac OS, Linux and FreeBSD, this permits writing according to the umask, like “`chmod +w`”.
- a** under Windows displays the file’s current attributes (Hidden, System, Read-only and Archive) and lets you specify a new attribute list. You can set the `dired-layout` variable under Windows to include these attributes in the dired listing itself. Under Unix, A runs the `chmod` command, passing it the mode

specification you type, such as `g+w` to let group members write to the file. For remote files accessed via `Scp`, Epsilon passes the mode specification to the remote system. It must be in the form of a Unix-style octal mode setting, like `0644`.

- r** refreshes the current listing. Epsilon will use the original file pattern to rebuild the file listing. If you've marked files for copying, moving, or deleting, the markings will be discarded if you refresh the listing, so Epsilon will prompt first to confirm that you want to do this.
- toggles whether Epsilon hides files and directories whose names that start with a `.` period character. See the `dired-show-dotfiles` variable to set which sorts of files this key hides.
- s** controls sorting. It prompts you to enter another letter to change the sorting method. Type `?` at that prompt to see the sorting options available.
- g** uses the directory associated with the current line to set Epsilon's current directory.
- +** creates a subdirectory. It asks for the new subdirectory's name.
- . or ^** invokes a `dired` on the parent directory of the current `dired`.
- 1** makes the window occupy the whole screen, then acts like `e`.
- 2 or 5** splits the window horizontally or vertically, then acts like `e` in the new window.
- o** switches to the next window, then acts like `e`.
- z** zooms the current window like the `zoom-window` command, then acts like `e`.
- !** prompts for a command line, then runs the specified program, adding the name of the current line's file after it. If the command line you type contains an `*`, Epsilon substitutes the current file name at that position instead of at the end. If the command line ends in a `&` character, Epsilon runs the program asynchronously; otherwise it waits for the program to finish.
- Shift-U or Shift-L** marks a file for uppercasing or lowercasing its file name, respectively. Press `X` to rename the marked files, as with other renaming keys. (Note that Epsilon for Windows displays all-uppercase file names in lowercase by default, so `Shift-U`'s effect may not be visible within Epsilon. See `preserve-filename-case`.)
- Shift-R** marks a file for a regular-expression replacement on its name. When you press `X` to execute operations on marked files, Epsilon will ask for a pattern and replacement text. Then for each marked file, it will perform the indicated replacement on its name to create a new file name, then rename the file to the new name. For instance, to rename a group of files like `dir\file1.cxx`, `dir\file2.cxx`, etc. to `dir2\file1.cpp`, `dir2\file2.cpp`, use `Shift-R` and specify `dir\(.*)\.cxx` as the search text and `dir2\#1.cpp` as the replacement text. To rename some `.htm` files to `.html`, specify `.*` as the search text and `#01` as the replacement text.
- Shift-G** marks files that contain some specific text. This subcommand prompts for some search text. You can use the keys `Ctrl-t`, `Ctrl-w` or `Ctrl-c` when typing the search string to toggle regex mode, word mode, or case folding. Then the subcommand prompts for a key to indicate what kind of marking to apply. Press `d`, `m`, or `c` to mark files for deletion, moving or copying, `u` to remove such markings, `U`, `L`, or `R` to perform the corresponding renaming function described above, or `g` to apply a generic marking that simply indicates which files contained the specified search string. A numeric prefix argument to this

subcommand reverses the sense of its test, marking only files that don't contain the specified text.

Alt-] moves forward to the next line with a different mark. For instance, if the current line indicates a file marked for deletion, it moves to the next line with a file that's not marked for deletion.

Alt-[moves back to the previous line with a different mark.

f runs the incremental-search command to search for text in a file name, restricting matches to the file name column.

Ctrl-c Alt-n copies the name of the file on the current line, including its full path, to the clipboard, so you can paste it elsewhere.

Shift-P prints the current file using the print-buffer command.

h or ? gives this help.

dired-sort	Dired mode: S	Sort a directory listing differently.
------------	---------------	---------------------------------------

In a dired buffer, this subcommand controls sorting. It prompts you to enter another letter to change the sorting method. Press N, E, S, or D to select sorting by file name, file extension, size, or time and date of modification, respectively. Press U to turn off sorting the next time Epsilon makes a dired listing, and display the file names in the same order they come from the operating system. (You can have Epsilon rebuild the current listing using the R dired subcommand.)

Press + or - at the sorting prompt to sort in ascending or descending order, respectively, or R to reverse the current sorting order. Press **<Enter>** to sort again using the currently selected sorting order.

Press G at the sorting prompt to toggle directory grouping. With directory grouping, Epsilon puts all subdirectories first in the list, then all files, and sorts each part individually. Without directory grouping, it mixes the two together (although it still puts . and . . first).

display-buffer-info	Brief: Alt-f	Display the name of the current file.
---------------------	--------------	---------------------------------------

This command displays the name of the file associated with the current buffer, and the mode of the current buffer. It displays an asterisk after the file name if the file has unsaved changes. This command can be useful if you've set Epsilon so it doesn't display these things continuously.

do-c-indent	C mode: <Tab>	Indent this line for C.
-------------	----------------------------	-------------------------

In a line's indentation, reindent the line correctly for C code. Inside the text of a line, or when repeated, insert a tab.

If a region is highlighted, Epsilon indents all lines in the region by one tab stop. With a numeric prefix argument, Epsilon indents by that amount.

do-vbasic-indent	<Tab> in VBasic mode	Indent this line for Visual Basic.
------------------	-----------------------------------	------------------------------------

In a line's indentation, reindent the line correctly for Tcl code. Inside the text of a line, or when repeated, insert a tab.

If a region is highlighted, Epsilon indents all lines in the region by one tab stop. With a numeric prefix argument, Epsilon indents by that amount.

down-line Ctrl-n Move point to the next line.

This command keeps point near the same horizontal position as it occupied on the previous line, if possible.

edit-customizations Load and edit your einit.ecm file to change settings.

This command locates and reads the einit.ecm file, in the same way as the find-file command. Epsilon reads and executes this file each time it starts.

edit-variables Alt-F8 Interactively set variables from a list.

This command displays a list of all variables and lets you set them. You can use the arrow keys or the normal movement keys to move around the list, or begin typing a variable name to have Epsilon jump to that portion of the list. Press <Enter> to set the highlighted variable, then edit the value shown using normal Epsilon commands.

To exit from edit-variables, press <Esc> or Ctrl-g.

With a numeric argument, the command includes system variables in its list.

In Epsilon for Windows, this command displays a list of variables. You can choose one, see its description and its current value, and modify it. The command will only list those variables included in the help file.

end-kbd-macro Ctrl-x) Stop defining a keyboard macro.

This command completes the keyboard macro started by the start-kbd-macro command. You may then execute the macro with the command last-kbd-macro, or you may give the macro a name with the commands bind-last-macro or name-kbd-macro.

end-of-line Ctrl-e Go to the end of the line.

This command positions point at the end of the current line, just before the newline character.

end-of-window Alt-. Go to last character in window.

Position point before the last character in the current window.

enlarge-window Ctrl-(PgUp) Enlarge window by one line.

If possible, the mode line of the window on top of the current window moves up. Otherwise, the current window's mode line moves down. This command has no effect if it would make any window smaller than two lines, counting the mode line.

enlarge-window-horizontally Alt-(PgDn) Enlarge window by one column.

If possible, the left boundary of the current window moves to the left by one character. Otherwise, the current window's right boundary moves to the right. This command has no effect if it would make any window smaller than one character wide.

`enlarge-window-interactively` `Ctrl-x +` Use arrow keys to resize a window.

This command lets you interactively change the size of the current window. After you invoke the command, use the arrow keys to point to a window border. The indicated border moves in a direction so as to make the current window larger. Keep pressing arrow keys to move window borders. To switch from enlarging to shrinking, press the minus key. Thereafter, the arrow keys cause the window border to move in a direction so as to shrink the window. When the window looks right, press `<Enter>` to leave the command.

`enter-key` `Ctrl-m` Insert a newline character.

This command acts like `normal-character` but inserts a newline character regardless of the key that invoked it. In overwrite mode, the `<Enter>` key simply moves to the beginning of the next line.

`epsilon-html-look-up` `F1 h` Look up a topic in the HTML Epsilon manual.

This command prompts for a topic, then displays a section of the Epsilon manual that refers to that topic, using a web browser.

`epsilon-info-look-up` `F1 f` Look up a topic in the Epsilon manual.

This command prompts for a topic, then displays a section of the Epsilon manual that refers to that topic, using Info mode. It's like using the `epsilon-manual-info` command followed by the `info-index` command. In EEL source code, the identifier at point becomes the default topic.

`epsilon-keyboard` Load a default keyboard, undoing keyboard changes.

This command restores Epsilon's original keyboard arrangement after running the `brief-keyboard` or `cua-keyboard` commands, which see. It restores a "canned" keyboard arrangement from the file `epsilon.kbd`, which must be on the path.

`epsilon-manual-html` Display the HTML-format version of the Epsilon manual.

This command displays the Epsilon manual's table of contents using a web browser.

`epsilon-manual-info` Display the Info-format version of the Epsilon manual.

This command enters Info mode and jumps to the top node of Epsilon's manual.

`epsilon-manual` Display Epsilon's manual.

This command makes Epsilon display its on-line manual in a web browser or in WinHelp, as appropriate. Also see `epsilon-manual-info`.

`eval` Compute and display the value of an expression.

This command prompts for an expression, then computes and displays its value using the integrated EEL compiler. The expression may have a numeric or string type. With a numeric prefix argument, this command inserts the result into the current buffer instead of displaying it. Also see the `execute-eel` command.

You can append a semicolon followed by a printf-style formatting character or sequence, and Epsilon will use that to display the result. For instance, adding `;x` displays the result in hexadecimal.

exchange-point-and-mark Ctrl-x Ctrl-x Swap point and mark.

Some commands such as `kill-region` and `copy-region` operate on the text between the point and the mark.

execute-eel Execute a line of EEL code.

This command prompts for an EEL statement, then executes it using the integrated EEL compiler. Also see the `eval` command.

exit Ctrl-x Ctrl-c Exit the editor.

If you haven't saved all your files, Epsilon will display a list using `bufed` and ask if you really want to exit. If you prefix this command with a numeric argument, however, Epsilon will simply exit and not ask you about any unsaved buffers. With a numeric argument, it also skips writing a session file, preserving your current font settings, and similar things. It returns the numeric argument as its exit code (instead of zero, returned for a normal exit).

Also see the `process-warn-on-exit` variable.

exit-level Ctrl-x Ctrl-z Exit the current recursive edit.

If you have entered a recursive edit (typically from `query-replace`), this command exits the recursive edit (bringing you back to the `replace`), otherwise it invokes `exit`.

exit-process Type "exit" to the concurrent process.

This command tries to make the currently executing concurrent process stop, by typing "exit" to it. A standard command processor exits when it receives this command.

export-colors Save color settings to an EEL source file.

The `export-colors` command constructs an EEL source file of color settings based on the current color settings. Use it to get a human-readable version of your color selections.

file-query-replace Shift-F7 Replace text in many files or buffers.

This command prompts for the text to search for and the replacement text. Then it prompts for a file name which may contain wildcards. The command then performs a `query-replace` on each file that matches the pattern, going to each occurrence of the search text, and asking whether or not to replace it.

If the `use-grep-ignore-file-variables` variable is nonzero, Epsilon skips over any file with an extension listed in `grep-ignore-file-extensions`; by default some binary file types are excluded, or those that match the `grep-ignore-file-basename`, `grep-ignore-file-pattern`, or `grep-ignore-file-types` variables, or those bigger than `grep-ignore-file-max-size`.

With a numeric argument, the command instead searches through all buffers. The buffer name pattern may contain the wildcard characters `?` to match any single character, `*` to match zero or more characters, or a character class like `[^a-zA-Z]` to match any non-alphabetic character.

At each occurrence of the search text, you have these choices:

y or `<Space>` replaces and goes to the next match.
n or `<Backspace>` doesn't replace, but goes to the next match.
`<Esc>` exits immediately.
. (`<Period>`) replaces and then exits.
`^` (`<Caret>`) backs up to the previous match, as long as it's within the same file.
! replaces all remaining occurrences in the current file without prompting, then asks if you want to replace all occurrences without prompting in all remaining files.
, (`<Comma>`) replaces the current match but doesn't go to the next match.
Ctrl-r enters a recursive edit, allowing you to modify the buffer arbitrarily. When you exit the recursive edit with `exit-level`, the replacement continues.
Ctrl-g exits and returns point to its original location in the current buffer, then asks if you want to look for possible replacements in the remaining files.
Ctrl-w toggles the state of word mode.
Ctrl-t toggles the state of regular expression mode (see the description of `regex-replace`).
Ctrl-c toggles the state of case-folding.
Any other key causes the command to skip to the next file.

The command doesn't save modified files back to disk. You can use the `save-all-buffers` command on `Ctrl-x s` to do this.

fill-comment	Various modes: Alt-q	Reformat the current paragraph in a comment.
--------------	----------------------	---

This command fills the current paragraph in a programming language comment, so that each line but the last becomes as long as possible without going past the fill column. It tries to preserve any prefix before each line. It uses language-specific patterns for recognizing comments, with special logic for C/C++/Java comments.

fill-indented-paragraph	Fill paragraph preserving indentation.
-------------------------	--

This command fills the current paragraph, so that each line but the last becomes as long as possible without going past the fill column. It tries to preserve any indentation before each line of the paragraph. In detail, it finds the line in the paragraph with the least amount of indentation (not counting the first line, whose indentation is always preserved), and preserves indentation to the left of that column, filling the text to the right of that column.

With a numeric argument, it fills the paragraph using the current column as the right margin, instead of the `margin-right` variable.

With a highlighted region, it fills each paragraph in the region.

fill-paragraph	Alt-q	Fill the current paragraph.
----------------	-------	-----------------------------

This command fills the current paragraph, so that each line but the last becomes as long as possible without going past the fill column. This command does not right-justify the paragraph with respect to the fill column.

With a numeric argument greater than 5, the paragraph is filled using that value as a temporary right margin. With a smaller numeric argument, the paragraph is filled using an infinite right margin, so all text goes on one long line.

With a highlighted region, it fills each paragraph in the region.

fill-region Fill the current region between point and mark.

This command fills each paragraph in the region between point and mark as in `fill-paragraph`. For this command, only completely empty lines separate one paragraph from another.

With a numeric argument greater than 5, the paragraph is filled using that value as a temporary right margin. With a smaller numeric argument, the paragraph is filled using an infinite right margin, so all text goes on one long line.

filter-region `Alt-|` Send the current region through an external program.

This command prompts for the name of a program and runs it, passing the current region to it as its standard input. It then displays any output from the program in a separate buffer. With a non-negative prefix argument, it replaces the current region with the program's output. With a negative prefix argument, it prompts for the name of a destination buffer. The command runs asynchronously except when it's used to replace the current region; see the `process-echo` variable to change that.

find-delimiter `Alt-)` Show the matching left delimiter.

This command shows the left parenthesis, square bracket, brace, or similar delimiter in a balanced expression. (Each mode defines its own list of delimiters used in that language.) It invokes `backward-level`, displays this location, pauses, and then returns point to its original location. Note that the cursor must appear after a right delimiter, not on it, to show the match for that delimiter.

You may change the length of time that this command pauses at the left delimiter by setting the variables `near-pause` and `far-pause`. The former specifies how long to pause (in hundredths of a second) if the left delimiter appeared in the window originally. The latter specifies how long to pause otherwise.

Regardless of the length of the pause, the pausing stops when you press a key.

find-file `Ctrl-x Ctrl-f` Put a file in the current window.

You would normally use this command to specify a file to edit. This command prompts you for a file name, then scans the buffers to see if any of them contain that file. If so, the command displays that buffer in the current window.

Otherwise, the command creates a buffer with the same name as the file, possibly modified to make it different from the names of nonempty buffers, then reads the file into this buffer, then displays that buffer in the current window.

Epsilon auto-detects the line termination convention of the file and performs any necessary translation. (See `set-line-translate`.) If the file uses a Unicode encoding, it detects that too, so long as the file begins with a byte order mark. (See `set-encoding`.) With a numeric argument, the command prompts for the desired translation and encoding methods.

If you simply type `(Enter)` for a file name, the command invokes `dired` with the current directory for the file pattern. Similarly, if you specify a directory or a file name with wild card characters, the command invokes `dired` with that pattern.

See the descriptions of the `prompt-with-buffer-directory` and `want-common-file-dialog` variables for more information on this command.

find-linked-file

Ctrl-x Ctrl-l

Grab the file name on this line and edit it.

Look on the current line for a file name, and edit that file like the `find-file` command. Epsilon uses special rules for certain modes. For HTML mode it looks for `href=` and `src=` attributes. For CSS it looks for `@import` directives. For C/C++/Java mode it follows `#include` references via the `include-directories` variable. In Java files it understands the `package` and `import` keywords, and looks along the `CLASSPATH` for packages. For files with a `.lst` extension, it assumes the current line holds a file name, instead of searching for a pattern that matches a typical file name. You can highlight a file name first if Epsilon has trouble picking it out.

find-oem-file

Read a file that uses the DOS character set.

Windows programs typically use a different character set than do DOS programs. The DOS character set is known as the DOS/OEM character set, and includes various line drawing characters and miscellaneous characters not in the Windows/ANSI set. The Windows/ANSI character set includes many accented characters not in the DOS/OEM character set. Epsilon for Windows uses the Windows/ANSI character set (with most fonts).

The `find-oem-file` command reads a file using the DOS/OEM character set, translating it into the Windows/ANSI character set, and arranges things so when you save the file, the reverse translation automatically occurs. This command is only available in Epsilon for Windows. See the `default-character-set` variable.

find-read-only-file

Edit a file preventing changes to it.

Prompt for a file name and edit the specified file, like the find-file command. Set the buffer read-only, and mark it so attempts to save the file prompt for a different name.

find-unconverted-file

Read a file without changing its character set.

If you've configured Epsilon for Windows to convert from the DOS/OEM character set to the ANSI character set upon reading a file, and to perform the opposite conversion when writing (by setting the `default-character-set` variable), use this command to bypass the conversion for a particular file.

follow-mode

Link this window with the next.

This command toggles follow mode, a buffer-specific minor mode. In follow mode, when the same buffer is displayed in two adjacent windows, moving and scrolling in one causes the other to scroll too, so that the two windows show adjoining sections of the buffer, with an overlap specified by the `follow-mode-overlap` variable. Follow mode can be convenient when editing narrow text on a wide display using side-by-side windows.

force-common-file-dialog

Make the next command use a common file dialog.

Commands in Epsilon for Windows that prompt for a file name can either use Epsilon's built-in prompt or a Windows common file dialog like other programs. The `want-common-file-dialog` variable controls this. By default, Epsilon uses its own dialog when you run such commands via key presses and the Windows common dialog when you run them using the mouse.

This command temporarily toggles this choice for the very next command. When bound to a key like Alt-F6, the key serves as a prefix key. So if typing Ctrl-x Ctrl-f normally has Epsilon use its own dialog, then Alt-F6 Ctrl-x Ctrl-f makes Epsilon use the Windows one, and vice versa.

forward-character Ctrl-f Go forward one character.

Nothing happens if you run this command with point at the end of the buffer. If you set `virtual-space` to two, the command can position the cursor past the last character on a line, and never moves to a different line.

forward-ifdef C mode: Alt-], Alt-⟨Down⟩ Find matching preprocessor line.

This command moves to the next `#if/#else/#endif` (or similar) preprocessor line. When starting from such a line, Epsilon finds the next matching one, skipping over inner nested preprocessor lines.

forward-level Ctrl-Alt-f Move point past a bracketed expression.

Point moves forward searching for one of `(`, `{`, or `[`. Then point moves past the nested expression. Point appears after the corresponding right delimiter. (Actually, each mode defines its own list of delimiters. One mode might only recognize `<` and `>` as delimiters, for instance.)

forward-paragraph Alt-] Go to the next paragraph.

Point travels forward through the buffer until it appears at the beginning of a paragraph. Blank lines (containing only spaces and tabs) always separate paragraphs.

You can control what Epsilon considers a paragraph using two variables.

If the buffer-specific variable `indents-separate-paragraphs` has a nonzero value, then a paragraph also begins with a nonblank line that starts with a tab or a space.

If the buffer-specific variable `tex-paragraphs` has a nonzero value, then Epsilon will not consider as part of a paragraph any sequence of lines that each start with at sign or period, if that sequence appears next to a blank line. And lines starting with `\begin` or `\end` or `%` will also delimit paragraphs.

forward-search-again Search forward for the same search string.

forward-sentence Alt-e Go to the end of the sentence.

Point travels forward through the buffer until positioned at the end of a sentence. A sentence ends with a period, exclamation point, or question mark, followed by two spaces or a newline, with any number of closing characters `"`, `'`, `)`, `]`, between. A sentence also ends at the end of a paragraph.

forward-word Alt-f Move past the next word.

By default, a word consists of a sequence of letters or underscores. The buffer-specific variable `word-pattern` contains a regular expression that defines Epsilon's notion of a word for the current buffer.

fundamental-mode

Turn off any special key definitions.

This command removes changes to key bindings made by modes such as C mode or Direed mode.

Every buffer has a major mode, and whenever you type keys in that buffer, Epsilon interprets them according to the buffer's mode. Each of Epsilon's various modes is suitable for editing a particular kind of text. Some modes only change the meanings of a few keys. For instance, C mode makes the `<Tab>` key indent the current line of C code. Other modes provide a group of new commands, usually on the letter keys. For example, in Direed mode the D key deletes a file. Each major mode is also the name of a command which puts the current buffer in that mode. For example, `Alt-x c-mode` puts the current buffer in C mode.

The default mode for new buffers you create with `select-buffer` is Fundamental Mode. (But see `new-file`.) This command returns the current buffer to Fundamental Mode, removing any changes to key bindings installed by another mode.

gams-mode

Set up for editing GAMS files.

This command sets up syntax highlighting suitable for files in the GAMS language used for mathematical programming.

<code>goto-beginning</code>	<code>Alt-<</code>	Go to the beginning of the buffer.
<code>goto-end</code>	<code>Alt-></code>	Go to the end of the buffer.
<code>goto-line</code>	<code>Ctrl-x g</code>	Go to a certain line by number.

This command moves point to the start of the n 'th line in the file, where n denotes the command's numeric argument. With no numeric argument, Epsilon will ask for the line number. You may add `:col` after the line number (or in place of it) to specify a column. The syntax `p12345` goes to the specified byte offset instead.

<code>goto-next-declaration</code>	<code>Ctrl-c Ctrl-n</code>	Move to next global definition or declaration.
------------------------------------	----------------------------	--

This command moves in source code files to the start of the next function definition, global variable definition, declaration of either of these, or any other item the current mode's tagging is able to tag.

<code>goto-next-definition</code>		Move to next global definition.
-----------------------------------	--	---------------------------------

This command moves in source code files to the start of the next function or global variable definition, or any other item the current mode's tagging is able to tag (but excluding declarations, if the current mode's tagger distinguishes these).

<code>goto-previous-declaration</code>	<code>Ctrl-c Ctrl-n</code>	Move to previous global definition or declaration.
--	----------------------------	--

This command moves in source code files to the start of the previous function definition, global variable definition, declaration of either of these, or any other item the current mode's tagging is able to tag.

goto-previous-definition

Move to previous global definition.

This command moves in source code files to the start of the previous function or global variable definition, or any other item the current mode's tagging is able to tag (but excluding declarations, if the current mode's tagger distinguishes these).

goto-tag

Ctrl-x .

Ask for the name of a function, then go there.

The command prompts you for the name of a tagged function, with completion. Epsilon then goes to the file and line where the function definition appears. If you give no name, Epsilon goes to the next tag in the alphabetical tag list. With a nonzero numeric argument, it goes to the next tag without asking for a tag name. Before moving to the tag, it sets a bookmark at the current position like `set-bookmark`.

A numeric argument of zero forces the definition to appear in a specific window. Epsilon prompts for a key to indicate which one. Press an arrow key to display the definition in the next window in that direction. Press `n` or `p` to display the definition in the next or previous window. Type the period character `.` to force the definition to appear in the current window. Press `2` or `5` to split the current window horizontally or vertically, respectively, and display the definition in the new window, or `1` to delete all windows but the current one, or `z` to run the `zoom-window` command first.

grep

Alt-F7

Search multiple files or buffers for a pattern.

This command lets you search a set of files for a pattern. It prompts for the search string and the file pattern. Then it scans the files, accumulating matching lines in the `grep` buffer. The `grep` buffer appears in the current window. By default, the `grep` command interprets the search string as a regular expression. Press `Ctrl-t` at the search string prompt to toggle regular expression mode. You can also type `Ctrl-w` or `Ctrl-c` to toggle word-mode or case-folding searches, respectively.

At the file pattern prompt, you can press `<Enter>` if you want Epsilon to search the same set of files as before. Type `Ctrl-s` and Epsilon will type in the directory part of the current buffer's file name; this is convenient when you want to search other files in the same directory as the current file. As at other prompts, you can also press `Alt-<Up>` key or `Alt-Ctrl-p` to show a list of your previous responses to the prompt. Use the arrow keys or the mouse to choose a previous response to repeat, and press `<Enter>`. If you want to edit the response first, press `Alt-e`.

You can use extended file patterns to search in multiple directories using a pattern like `**.{c,cpp,h}` (which searches in the current directory tree for `.c`, `.cpp`, and `.h` files).

If the `use-grep-ignore-file-variables` variable is nonzero, Epsilon skips over any file with an extension listed in `grep-ignore-file-extensions`; by default some binary file types are excluded, or those that match the `grep-ignore-file-basename`, `grep-ignore-file-pattern`, or `grep-ignore-file-types` variables, or those bigger than `grep-ignore-file-max-size`.

With a numeric argument, `grep` instead searches through buffers, defaulting to the current buffer. The buffer name pattern may contain the wildcard characters `?` to match any single character, `*` to match zero or more characters, a character class like `[^a-zA-Z]` to match any non-alphabetic character, or `|` to separate alternatives. (The `buffer-grep` command provides this functionality as a separate command.)

In `grep` mode, alphabetic keys run special `grep` commands. See the description of the `grep-mode` command for details. Typing `H` or `'?`' in `grep` mode gives help on `grep` subcommands.

grep-mode

Edit a list of lines containing a search string.

In a grep buffer, you can move around by using the normal movement commands. Most alphabetic keys run special grep commands. The ‘N’ and ‘P’ keys move to the next and previous matches; Alt-n and Alt-p move to the next and previous files. You can easily go from the grep buffer to the corresponding locations in the original files. To do this, simply position point on the copy of the line, then press (Space), (Enter), or ‘E’. The file appears in the current window, with point positioned at the beginning of the matching line. Typing ‘1’ brings up the file in a window that occupies the entire screen. Typing ‘2’ splits the window horizontally, then brings up the file in the lower window. Typing ‘5’ splits the window vertically, then brings up the file. Typing the letter ‘O’ shows the file in the next window on the screen, without splitting windows any further. Typing ‘Z’ runs the zoom-window command, then brings up the file.

help

F1

Get documentation on commands.

If executed during another command, help simply pops up the description of that command. Otherwise, you press another key to specify one of the following options:

- ?** prints out this message.
- k** runs `describe-key`, which asks for the key, then gives full help on the command bound to that key.
- c** runs `describe-command`, which asks for the command name, then gives full help on that command, along with its bindings.
- r** runs `describe-variable`, which asks for the variable name, then shows the full help on that variable.
- i** runs the `info` command, which starts Info mode. Info mode lets you read the entire Epsilon manual, as well as any other documentation you may have in Info format.
- Ctrl-c** runs the `info-goto-epsilon-command` command, which prompts for the name of an Epsilon command, then displays an Info page from Epsilon’s online manual that describes the command.
- Ctrl-k** runs the `info-goto-epsilon-key` command, which prompts for a key, then displays an Info page from Epsilon’s online manual that describes the command it runs.
- Ctrl-v** runs the `info-goto-epsilon-variable` command, which prompts for an Epsilon variable’s name, then displays an Info page from Epsilon’s online manual that describes that variable.
- f** runs the `epsilon-info-look-up` command, which prompts for a topic, then starts Info mode and looks up that topic in the Epsilon manual.
- h** displays Epsilon’s manual in HTML format, by running a web browser. It prompts for a topic, which can be a command or variable name, or any other text. (The browser will try to find an exact match for what you type; if not, it will search for web pages containing that word.) When you’re looking at Epsilon’s manual in Info mode, using one of the previous commands, this command will default to showing the same topic in a browser.
- w** runs the WinHelp program to display Epsilon’s online manual, in Epsilon for Windows versions prior to Windows Vista.
- a** runs `apropos` which asks for a string, then lists commands and variables apropos that string.
- b** runs `show-bindings`, which asks for a command name, then gives you its bindings.

? shows help on hex mode.

highlight-region Ctrl-x Ctrl-h Highlight area between point and mark.

This command toggles highlighting of the region of the buffer between point and mark. If you prefix a nonzero numeric argument, the command highlights the region; a numeric argument of zero turns highlighting off.

html-backward-tag HTML mode: Alt-Shift-B Move backward over a tag or element.

When point is at an HTML or XML tag, this command moves to its start. At an end tag that has a matching start tag, it moves before the start tag. It moves by words when outside a tag.

html-close-last-tag HTML mode: Alt-Shift-E Insert an end tag.

This command looks back for the most recent unclosed HTML or XML start tag, then inserts an end tag to close it.

html-delete-tag HTML mode: Alt-Shift-D Delete the current tag and its mate.

This command deletes the HTML or XML tag at point (or just after point). If the tag is a start tag that has a matching end tag, or an end tag with a matching start tag, the command deletes the matching tag also.

html-fill-paragraph HTML mode: Alt-q Fill the current paragraph.

This command fills the current paragraph using the fill-paragraph command. But within embedded scripting, it calls the fill-comment command instead.

html-find-matching-tag HTML mode: Alt-= Go to tag matching this one.

In HTML and XML modes, this command moves to the end tag that matches the start tag at point, or vice versa.

html-forward-tag HTML mode: Alt-Shift-F Move forward over a tag or element.

When point is at an HTML or XML tag, this command moves to its end. At a start tag that has a matching end tag, it moves after the end tag. It moves by words when outside a tag.

html-indent-cmd ⌵Tab⌶ in HTML mode Indent this line for HTML or XML.

In a line's indentation, reindent the line correctly for HTML or XML code. Inside the text of a line, or when repeated, insert a tab.

If a region is highlighted, Epsilon indents all lines in the region by one tab stop. With a numeric prefix argument, Epsilon indents by that amount.

html-list-element-nesting HTML mode: Alt-i Show element nesting in effect at point.

This command displays a dialog showing the nesting of HTML or XML elements that apply at point. It indicates point is inside an open tag by showing <tagname instead of <tagname>. Unclosed elements and some other basic syntax problems are also reported.

Unlike Epsilon's continuous nesting display on the mode line (see the `html-display-definition` variable), which shows nesting at the beginning of the current line, and doesn't change as you move around within a line, this command uses the current position within the line.

html-list-mismatched-tags HTML mode: Alt-Shift-L Show unpaired tags.

This command finds all start or end tags in the current buffer without a match. It lists them into a buffer in grep-mode, which see. If HTML permits a tag to be omitted, it won't be listed.

html-mode Set up for editing Hypertext Markup Language files.

This command puts the current buffer in HTML mode. Epsilon will do syntax-highlighting, smart indenting, and brace-matching.

html-redirect-active-key various Self-insert, fix indentation.

Delimiter keys bound to this command in HTML/XML modes self-insert, optionally display their matching delimiter, then reindent as appropriate.

import-colors Convert color choices to EEL source code.

This command constructs an EEL source file with color settings based on a "changes" file produced when updating from an older version of Epsilon. It may be helpful if you prefer to keep color settings in an EEL file.

import-customizations Build a list of prior customizations.

This command constructs an einit.ecm customization file based on the customizations recorded in your state file for a previous version of Epsilon. It looks for variable definitions, key bindings, color settings and so forth. The new definitions are added to the end of your einit.ecm file, and Epsilon will use them the next time it starts. Epsilon comments out any previous settings in that file, and deletes previous settings it commented out on a previous run. With a numeric argument, it doesn't comment out or delete previous settings.

It works by running the previous version of Epsilon with no customizations loaded, having it generate a complete list of settings, and then comparing that list to the results when the previous version of Epsilon runs with your customized settings. Each difference represents a particular customization.

First it displays the settings it will use for the above process. You can edit them if necessary, for instance to refer to the files of a different installed version that Epsilon didn't automatically locate.

If your previous configuration loaded any EEL extension files to define new commands, you should list them on the settings page shown.

Also see the clean-customizations and list-customizations commands.

incremental-search Ctrl-s Search for a string as you type it.

Ctrl-q quotes the next character. Backspace cancels the last character. Ctrl-s repeats a forward search, and Ctrl-r repeats a backward search, or they change its direction. Ctrl-r or Ctrl-s with an empty search string brings back the search string from the previous search. Ctrl-o enables or disables incremental mode. Incremental mode searches as you type; non-incremental mode lets you edit the search string.

Ctrl-w enables or disables word searching, restricting matches to complete words. Ctrl-t enables or disables regular expression searching, in which the search string specifies a pattern

(see `regex-search` for rules). `Ctrl-c` enables or disables case-folding. `<Enter>` or `<Esc>` exits the search, leaving point alone.

If Epsilon cannot find all the input string, it doesn't discard the portion it cannot find. You can delete it, discard it all with `Ctrl-g`, use `Ctrl-r` or `Ctrl-s` to search the other way, change modes, or exit from the search.

Press `Alt-(Up)` or `Ctrl-Alt-p` to select from a list of previous search strings. Press `Alt-g` to retrieve the last search string for the current buffer.

During incremental searching, if you type Control or Alt keys not mentioned above, Epsilon exits the search and executes the command bound to the key. During a non-incremental search, most Control and Alt keys edit the search string itself.

Quitting (with `Ctrl-g`) a successful search aborts the search and moves point back; quitting a failing search just discards the portion of the search string that Epsilon could not find.

indent-for-comment	Alt-;	Indent and insert a comment, or search for one.
--------------------	-------	---

This command creates a comment on the current line, using the commenting style of the current language mode. The comment begins at the column specified by the `comment-column` variable (by default 40). (However, if the comment is the first thing on the line and `indent-comment-as-code` is nonzero, it indents to the column specified by the buffer's language-specific indentation function.) If the line already has a comment, this command reindents the comment to the comment column.

With a numeric argument, this command doesn't insert a comment, but instead searches for one. With a negative numeric argument, it searches backwards for a comment.

indent-previous	<Tab>	Indent based on the previous line.
-----------------	-------	------------------------------------

This command makes the current line start at the same column as the previous non-blank line. Specifically, if you invoke this command with point in or adjacent to a line's indentation, `indent-previous` replaces that indentation with the indentation of the previous non-blank line. If point's indentation exceeds that of the previous non-blank line, or if you invoke this command with point outside of the line's indentation, this command simply runs `indent-to-tab-stop`.

If a region is highlighted, Epsilon indents all lines in the region by one tab stop. With a numeric prefix argument, Epsilon indents by that amount.

indent-region	Ctrl-Alt-\	Indent from point to mark using the function on <Tab>.
---------------	------------	--

This command goes to the start of each line in the region and does what the `<Tab>` key would do if pressed. It then deletes any resulting lines that contain only spaces and tabs, replacing them with newline characters.

indent-rigidly	Ctrl-x Ctrl-i	Move all lines in the region left or right by a fixed amount.
----------------	---------------	---

This command finds the indentation of each line in the region, and augments it by the value of the numeric argument. With a negative numeric argument, $-n$, the command removes n columns from each line's indentation.

With no numeric argument it uses the variable `soft-tab-size` if it's nonzero. Otherwise it uses `tab-size`.

You can also invoke this command by highlighting the region and pressing `<Tab>` or `Shift-<Tab>` to add or subtract indentation.

`indent-to-tab-stop` Indent to the next tab stop.

This command inserts spaces (and perhaps tabs, depending on the `indent-with-tabs` variable) to reach the next tab stop.

It checks the variable `soft-tab-size`, the current `thismode-soft-tab-size` variable, if any, where `thismode` is the current mode's name, and the `tab-size` variable, and uses the first of these that's greater than zero.

If a region is highlighted, Epsilon indents all lines in the region by one tab stop. With a numeric prefix argument, Epsilon indents by that amount.

`indent-under` Ctrl-Alt-i Indent to the next text on the previous line.

This function starts at the current column on the previous non-blank line, and moves right until it reaches the next column where a run of non-space characters starts. It then replaces the indentation at point with indentation that reaches to the same column, by inserting tabs and spaces. With a highlighted region, it indents all lines in the region to that same column, scanning the last non-blank line before the region.

With a numeric prefix argument, `indent-under` goes to a different run of non-spaces. For instance, with an argument of 3, it goes to the previous line and finds the third word after the original column, then aligns the original line there.

If the previous non-blank line doesn't have the requested number of words, it runs `indent-to-tab-stop`.

`info` F1 i Read documentation in Info format.

This command starts Epsilon's Info mode for reading Info-format documentation. Use 'q' to switch back to the previous buffer. Commands like `<Space>` and `<Backspace>`, N and P, navigate through the tree-structured Info hierarchy. See `info-mode` for details.

`info-backward-node` Info: [Walk the leaves of the Info hierarchy in reverse.

This command goes to the previous node in the sequence of Info nodes formed by walking the leaves of the hierarchy within the current Info file.

In detail, it goes to the previous node, then as long as it's on a node with a menu, goes to the last menu item. However, if there's no previous node (or it's the same as the current node's parent), it goes up to the parent node as long as it's in the same file.

`info-directory-node` Info: D Go to the Directory node.

Info nodes are arranged in a hierarchy. At the top of the hierarchy is one special node that contains links to each of the other Info files in the tree. This command goes to that topmost node.

`info-follow-nearest-reference` Info: `<Enter>` Follow the link near point.

After navigating among the cross references or menu items in an Info node with `<Tab>` or `<Backtab>` (or in any other way), use this key to follow the selected link.

info-follow-reference	Info: F	Prompt for a cross-reference in this node, then go there.
-----------------------	---------	---

This command prompts for the name of a cross-reference in this node, with completion, then goes to the selected node.

info-forward-node	Info:]	Walk the leaves of the Info hierarchy.
-------------------	---------	--

This command goes to the next node in the sequence of Info nodes formed by walking the leaves of the hierarchy within the current Info file.

In detail, if a menu is visible in the window, go to its next item after point. Otherwise, go to this node's next node. (If there is no next node, go up until reaching a node with a next node first, but never to the Top node.)

info-goto	Info: G	Ask for a node's name, then go there.
-----------	---------	---------------------------------------

This command prompts for the name of a node, then goes to it. It offers completion on the names of all the nodes in the current file, but you may also refer to a different file using a node name like (FileName)NodeName.

info-goto-epsilon-command	F1 Ctrl-c	Prompt for a command, look up Info.
---------------------------	-----------	-------------------------------------

This command prompts for the name of an Epsilon command, then displays an Info page from Epsilon's online manual that describes the command.

info-goto-epsilon-key	F1 Ctrl-k	Prompt for a key, look up Info.
-----------------------	-----------	---------------------------------

This command prompts for a key, then displays an Info page from Epsilon's online manual that describes the command it runs.

info-goto-epsilon-variable	F1 Ctrl-v	Prompt for a variable, look up Info.
----------------------------	-----------	--------------------------------------

This command prompts for an Epsilon variable's name, then displays an Info page from Epsilon's online manual that describes that variable.

info-index	Info: I	Prompt for an index entry; then go to its first reference.
------------	---------	--

This command prompts for some text, then goes to the destination of the first index entry containing that text. Use the info-index-next command on <Comma> to see other entries. If you just press <Enter> at the prompt, Epsilon goes to the first index node in the current Info file, and you can peruse the index entries yourself.

info-index-next	Info: <Comma>	Go to the next matching index entry.
-----------------	---------------	--------------------------------------

This command goes to the next index entry that matches the text specified by the most recent info-index command. Upon reaching the last item, it wraps and goes to the first matching item again.

info-last Info: L Return to the most recently visited node.

Info remembers the history of all nodes you've visited. This command goes to the last node on that list. Repeat it to revisit older and older nodes.

info-last-node Info: > Go to the last node in this file.

This command goes to the last node in this Info file. In detail, Epsilon goes to the top node of the file, goes to the last node in its menu, then follows Next nodes until there are no more, then moves like info-forward-node until it can move no further.

info-menu Info: M Prompt for a menu item, then go there.

This command prompts for the name of a menu item in this node's menu, with completion, then goes to the selected node.

info-mode Put this buffer in Info mode.

This command sets up keys for browsing an Info file. Normally you would run the info command, not this one.

These are the commands in Info mode:

- h** shows detailed documentation on using Info mode.
- ?** displays this list of available Info commands.
- <Space>** pages through the entire Info file one screenful at a time, scrolling either or moving to a different node as appropriate.
- <Backspace>** pages backwards through the Info file.
- <Tab>** moves to the next reference or menu item in this node.
- <Backtab>** moves to the previous reference or menu item in this node.
- <Enter>** follows the current reference or menu item to another node. You can also double-click one of these with the mouse to follow it.
- b** moves to the beginning of the current node.
- l** (letter L) goes to the most recently visited node before this one in the history list.
- n** goes to the next node after this one, as designated in the heading at the top of this node.
- p** goes to the previous node before this one, as designated in the heading at the top of this node.
- u** goes up to the parent of this node, as designated in the heading at the top of this node.
- m** prompts for the name of an entry in this node's menu, then goes to it.
- 1, 2, 3, ... 0** goes to the first, second, third, ... entry in this node's menu. 0 goes to the last entry in this node's menu.
- f** prompts for the name of a cross-reference in this node, then goes to it.
- t** goes to the top node in the current Info file, which is always named Top.
- d** goes to the directory node, a node that refers to all known Info files. From here you can navigate to any other Info file.
- g** prompts for the name of a node, then goes to it.

-]
] goes to the next node in the sequence of Info nodes formed by walking the leaves of the hierarchy within the current Info file, much like `<Space>` but without paging.
- [
[goes to the previous node in the sequence of Info nodes formed by walking the leaves of the hierarchy within the current Info file, much like `<Backspace>` but without paging.
- >
> goes to the last node in the file, viewed as a hierarchy (the node a repeated] would eventually reach).
- s
s prompts for a search string, then searches for the next match, switching nodes if necessary. Keys like `Ctrl-t` to toggle regular expression mode work as usual. Use `Ctrl-s` or `Ctrl-r` instead of `S` to search only within the current node.
- i
i prompts for text, then looks it up in this Info file's indexes, and goes to the first node with an index entry containing that text. Press `<Enter>` without typing any text to just go to the first index.
- ,
, goes to the next entry in the set of index entries set by the last `I` command.
- q
q quits Info mode by switching this window to the buffer it displayed before you entered Info mode.

info-mouse-double Follow the selected link.

Double-clicking a link (a menu item in an Info node, a cross-reference, or the Next, Prev, or Up links at the top of a node) runs this command, which simply follows the link.

info-next Info: N Go to the next node after this one.

This command goes to the next node after this one, named in the current node's header line.

info-next-page Info: `<Space>` Page down, then move to the next node.

Use this command to page through the entire Info file one screenful at a time.

In detail, if a menu is visible in the window, this command goes to its next item after point. Otherwise, it tries to scroll down. Otherwise, it goes to this node's next node, going up the tree if necessary to find a node with a next node.

info-next-reference Info: `<Tab>` Move to the next reference or menu item.

This command moves forward to the next link in this node: either a reference or a menu item. Use `<Tab>` and `<Backtab>` to select a link, then `<Enter>` to follow it.

info-nth-menu-item Info: 1, 2, ..., 0 Follow that menu entry.

This command goes to a menu entry without prompting as `M` does. 1 goes to the first item in the menu, 2 to the second and so forth. 0 goes to the last item in the menu.

info-previous Info: P Go to the previous node before this one.

This command goes to the previous node before this one, named in the current node's header line.

info-previous-page Info: `<Backspace>` Page up, or move to a previous node.

Use this command to page backward through the entire Info file one screenful at a time.

In detail, if a menu is above point, go to its closest item and then keep following the last item in the current node's menu until reaching one without a menu. Otherwise (if the current node has no menu above point), page up if possible. Otherwise move to this node's previous node, and then keep following the last item in the current node's menu until reaching one without a menu. Otherwise (if the original node had no previous node, or its previous node was the same as its up node), move to the original node's up node (but never to a different file).

info-previous-reference Info: `<Backtab>` Move to the previous reference
or menu item.

This command moves backward to the previous link in this node: either a reference or a menu item. Use `<Tab>` and `<Backtab>` to select a link, then `<Enter>` to follow it.

info-quit Info: `Q` Exit Info mode.

This command leaves Info mode by switching this window to the buffer it displayed before you entered Info mode.

info-search Info: `S` Search for text in many nodes.

This command prompts for search text, then searches for the text, switching nodes if necessary. Keys like `Ctrl-t` to toggle regular expression mode work as usual. Use `Ctrl-s` or `Ctrl-r` instead of `S` to search only within the current node.

info-tagify Rebuild the tag table for this Info file.

Epsilon can more quickly navigate between the nodes of a big Info file if it has an up-to-date tag table. This command builds (or rebuilds) a tag table for the current Info file, and is useful after you edit an Info file. The tag table is stored in a special hidden node.

info-top Info: `T` Go to the top node in the current file.

This command goes to the top node in the current Info file, which is always named Top.

info-up Info: `U` Go to parent of this node.

This command goes to the parent of the current node, indicated with "Up:" in this node's header line.

info-validate Check an Info file for errors.

This command checks an Info file for certain common errors. It reports on menu items or cross-references that refer to non-existent nodes.

ini-mode A mode for editing .ini files.

This mode provides syntax highlighting suitable for MS-Windows .ini files.

jump-to-column Alt-g Go to the specified column.

This command prompts for a number, then moves to the specified column on the current line. In horizontal scrolling mode, it then horizontally centers the window on that column (or, if possible, positions the window so that the start of the line is also visible). You can specify the column with a numeric prefix argument and avoid the prompt.

jump-to-dvi TeX mode: Alt-Shift-J Show DVI or PDF output from this TeX line.

In a TeX buffer, this command tells a running DVI previewer or PDF viewer to display the output resulting from the text near point. See page 98 for instructions on setting this up. You can also left-click while holding down the Ctrl key.

jump-to-last-bookmark Alt-j Go to a previously recorded place.

Use this command to jump to a location that you previously set with the `set-bookmark` command. If you repeatedly press this key, you will cycle through the last 10 temporary bookmarks.

jump-to-named-bookmark Ctrl-x j Go to a named bookmark.

Use this command to jump to a location that you previously saved with the `set-named-bookmark` command. The command prompts for a bookmark name (a letter), then jumps to that bookmark.

If you specify a digit instead of a letter, the command jumps to the corresponding temporary bookmark (set with `set-bookmark`). Zero refers to the last such temporary bookmark, one to the previous one, and so on.

You can press ‘?’ to get a list of the currently defined bookmarks, along with the text that contains the bookmarks. To select one, simply move to the desired bookmark and press `(Enter)`.

keep-duplicate-lines Remove unduplicated lines.

This command deletes all lines that only occur once, and leaves one copy of each duplicated line. If the `case-fold` variable is nonzero, lines that only differ by case will be considered identical. Also see the `uniq` and `keep-unique-lines` command.

keep-matching-lines Delete all lines but those containing a regex pattern.

This command prompts for a regular expression pattern. It then deletes all lines below point in the current buffer except those that contain the pattern. While you type the pattern, Ctrl-w enables or disables word searching, restricting matches to complete words. Ctrl-t enables or disables regular expression searching, in which the search string specifies a pattern (see `regex-search` for rules). Ctrl-c enables or disables case-folding.

keep-unique-lines Entirely remove duplicate lines.

This command deletes all copies of any duplicated adjacent lines, leaving only those lines that match neither the preceding nor the following line. If the `case-fold` variable is nonzero, lines that only differ by case will be considered identical. Also see the `uniq` and `keep-duplicate-lines` command.

`kill-all-buffers` Delete all user buffers.

This command discards all of Epsilon's buffers (except hidden system buffers).

`kill-buffer` Ctrl-x k Make a specified buffer not exist.

This command asks for a buffer name and then deletes that buffer. The command warns you before deleting a buffer that contains unsaved changes.

`kill-comment` Kill the next comment.

This command searches forward for a comment, as defined by the current mode, then kills it. The `set-comment-column` command invokes this command if given a negative numeric argument.

`kill-current-buffer` Ctrl-x Ctrl-k Make the current buffer not exist.

This command deletes the current buffer and switches to another, creating a new buffer if necessary. The command warns you first if the current buffer contains unsaved changes.

`kill-current-line` Kill the current line.

This command kills the entire current line, including any newline at its end. The killed text goes to a kill buffer for possible later retrieval.

`kill-level` Ctrl-Alt-k Kill a bracketed expression.

The command moves point as in `forward-level`, killing the characters it passes over.

`kill-line` Ctrl-k Kill to end of line.

If invoked with point at the end of a line, this command kills the newline. Otherwise, it kills the rest of the line but not the newline. If you give `kill-line` a numeric argument, it kills that many lines and newlines. The killed text goes to a kill buffer for possible later retrieval.

`kill-process` Get rid of a concurrent process.

This command disconnects Epsilon from a concurrent process in the current buffer and makes it exit. If the current buffer isn't a process buffer but has a running Internet job (such as an `ftp://` buffer), this command tries to cancel it. If the current buffer isn't process buffer and doesn't have a running Internet job, the command tries to kill a process running in the buffer named "process".

`kill-rectangle` Kill the rectangular area between point and mark.

This command removes the characters in the rectangular area between point and mark, and puts them in a kill buffer. By default, the deleted area is replaced with spaces and tabs, and text to the right of the rectangle remains in the same position. With a numeric argument, this command removes the deleted rectangle and shifts leftward any text to its right. See the `kill-rectangle-removes` variable to make this command remove text by default. Also see the `delete-rectangle` command.

kill-region Ctrl-w Kill the text between point and mark.

This command removes the characters between point and mark from the buffer, and puts them in a kill buffer. For rectangular regions, it simply calls kill-rectangle.

kill-sentence Alt-k Kill to the end of the sentence.

The command moves point as in forward-sentence, killing the characters it passes over.

kill-to-end-of-line Brief: Alt-k Kill the remainder of the current line.

This command kills the remainder of the current line, not including any newline at its end. If point is at the end of the line, the command does nothing. The killed text goes to a kill buffer for possible later retrieval.

kill-window Ctrl-x 0 Delete the current window.

This command gets rid of the current window, and gives the space to some other window. This command does not delete the buffer displayed in the window.

kill-word Alt-d Kill the word after point.

The command moves point forward through the buffer as with forward-word, then kills the region it traversed.

last-kbd-macro Ctrl-F4 Execute the last keyboard macro defined from the keyboard.

This command runs the last keyboard macro you defined with the start-kbd-macro and end-kbd-macro commands.

latex-mode Set up for editing LaTeX documents.

This command sets up Epsilon for editing LaTeX documents. Keys in LaTeX mode include Alt-i for italic text, Alt-Shift-I for slanted text, Alt-Shift-T for typewriter, Alt-Shift-B for boldface, Alt-Shift-C for small caps, Alt-Shift-F for a footnote, and Alt-s for a centered line. Alt-Shift-E prompts for the name of a LaTeX environment, then inserts `\begin{env}` and `\end{env}` lines.

For all these commands, you can highlight a block of text first and Epsilon will make the text italic, slanted, etc. or you can use the command and then type the text to be italic, slanted, etc.

The keys ‘{’ and ‘\$’ insert matched pairs of characters (either {} or \$\$), the keys ⟨Comma⟩ and ⟨Period⟩ remove a preceding italic correction \/, the " key inserts the appropriate kind of doublequote sequence like ‘ ‘ or ’ ’, and Alt-" inserts an actual " character.

line-to-bottom Brief: Ctrl-b Scroll window to move this line to bottom.

This command tries to scroll the current window so that the line containing point becomes the last line in the window.

line-to-top Brief: Ctrl-t Scroll the window to move this line to the top.

This command tries to scroll the current window so that the line containing point becomes the first line in the window.

list-all Describe Epsilon's state in text form.

This command puts a description of Epsilon's state, including bindings, macros, variables, and commands, in a buffer named list-all. It provides complete descriptions for bindings, macros, and simple variables, but for commands and subroutines, it only records the fact that a function with that name exists. You can copy key binding or macro lines from this listing, modify them, and put them in your einit.ecm file to customize Epsilon, or load them using load-buffer.

list-bookmarks Pop up a list of all the bookmarks.

This command works like jump-to-named-bookmark, but pops up a list of bookmarks, as if you had typed '?' to that command. If you always want the pop up list, you can bind this command to a key (perhaps replacing the default binding of jump-to-named-bookmark on Ctrl-x j).

list-debug Show EEL functions set for debugging.

This command displays the names of all functions that have had debugging enabled using the set-debug command.

list-changes List variables added or changed when updating.

You would use this command when updating to a new version of Epsilon. It asks for the names of two files, then makes a list of all lines from the second that don't appear in the first. It sorts the second file, but not the first.

list-colors Make a list of all color settings.

This command constructs a buffer with all of Epsilon's current color settings, one to a line. The export-colors command is another way to save color selections in human-readable form.

list-customizations Build a list of current customizations.

This command constructs an einit.ecm customization file based on your current set of variable definitions, key bindings, color settings, loaded EEL extensions, and so forth. The new definitions are added to the end of your einit.ecm file, and Epsilon will use them the next time it starts. Epsilon comments out any previous settings in that file, and deletes previous settings it commented out on a previous run. With a numeric argument, it doesn't comment out or delete previous settings.

It works by running another copy of Epsilon with no customizations loaded, having it generate a complete list of settings, and then comparing that list to current settings. Each difference represents a particular customization.

Some definitions and settings may be the result of loading an EEL extension file. Rather than list these individually, this command lists just the EEL source file name. It displays all EEL source files it knows of, before it starts, and you can edit this list if you like.

Also see the clean-customizations and import-customizations commands.

list-definitions Alt-' List functions defined in this file.

This command displays a list of all functions and global variables defined in the current file. It uses Epsilon's tagging facility, so it works for any file type where tagging works.

You can move to a definition in the list and press <Enter> and Epsilon will go to that definition. Or press Ctrl-g to remain at the starting point.

By default, it skips over external declarations. With a prefix numeric argument, it includes those too. (If the buffer contains only external declarations and no definitions, a prefix argument is unnecessary; Epsilon will automatically include them.)

list-files Create a buffer listing all files matching a pattern.

This command prompts for a file name pattern containing wildcards, then creates a list of all the files matching the pattern in a buffer named "file-list". Use this command when you need a plain list of file names, without any of the extra information that the similar dired command provides.

With a numeric argument, the command lists matching directory names, as well as file names.

list-make-preprocessor-conditionals Makefile mode: Alt-i Show conditionals in effect for this line.

In makefile mode buffers, this command displays a list of all preprocessor conditionals that affect the current line.

list-preprocessor-conditionals C mode: Alt-i Show conditionals in effect for this line.

In C mode buffers, this command displays a list of all preprocessor conditionals that affect the current line.

list-undefined Which EEL functions are not defined anywhere?

This command makes a list of all EEL functions that are called from some other EEL function, but have not been defined. Epsilon doesn't report any error when you load an EEL function that refers to an undefined function, but you'll get an error message when the function runs. This command helps to prevent such errors. The list also includes any variables or functions that have been deleted.

load-buffer Interpret a buffer as a command file.

This command prompts you for the name of a buffer containing macro definitions and key bindings in command file format, then executes the commands contained in that buffer. For information on command file format, see the section of the manual entitled "Command Files".

load-bytes Load compiled EEL commands and variables.

This command prompts you for the name of a file produced by the EEL compiler, then loads that file. You may omit the file name's extension. The command changes any file name extension you provide to ".b".

`load-changes` Load the changes into Epsilon.

The `load-changes` command prompts for a file name, then loads the changes described in that file. Use this command when updating to a new version of Epsilon, to load the output of the `list-changes` command.

`load-file` Read in a command file.

This command prompts you for the name of a command file containing macro definitions and key bindings, then executes the commands contained in that file. For information on command file format, see the section of the manual entitled “Command Files”.

`locate-file` Search for a file.

This command prompts you for a file name and then searches for that file. In Windows, it searches for the file on all local hard drives, skipping over removable drives, CD-ROM drives, and network drives. On Unix, it searches through particular parts of the directory hierarchy specified by the `locate-path-unix` variable.

`lowercase-word` Alt-l Make the current word lower case.

Point travels forward through the buffer as with `forward-word`. It turns all the letters it encounters to lower case. If the current buffer contains a highlighted region, Epsilon instead changes all the letters in the region to lower case, leaving point unchanged.

`mail-fill-paragraph` Ctrl-c Alt-q Fill a paragraph preserving email quoting.

Email messages commonly use lines beginning with `>` or `#` to indicate quoting. This command fills paragraphs in such emails, preserving such quoting, and recognizing that different amounts of quoting serve to separate paragraphs. It uses the `mail-quote-pattern` and `mail-quote-skip` variables to determine the permissible quoting characters.

With a numeric argument, it fills paragraphs using the current column as the right margin, instead of the `margin-right` variable.

The `prefix-fill-paragraph` command is similar, but uses different rules that aren’t as well-adapted for email.

With a highlighted region, it fills each paragraph in the region.

`mail-quote-region` Ctrl-c > Quote the current region or paragraph for email.

This command inserts a `>` character before each line in the current paragraph, quoting it for email. When a region is highlighted, it instead inserts a `>` character before each line in the region. The actual text inserted comes from the `mail-quote-text` variable.

`mail-unquote` Ctrl-c < Remove quoting from the current quoted section.

This command removes email-style quoting. It examines the lines around point to find a range of lines with matching quoting, then removes the quoting (by default, `>` or `#` characters) from those lines. If a region has been highlighted, it removes quoting only from those lines. It uses the `mail-quote-pattern` and `mail-quote-skip` variables to determine the permissible quoting characters.

mark-normal-region Brief: Alt-a Begin marking a normal region.

This command begins marking and highlighting a region of text, defining it as a normal (non-inclusive) region. A normal region includes all the characters between point and mark. When you run this command, it sets the mark equal to the value of point, so initially the highlighted region is empty.

If Epsilon is already highlighting a region of another type, this command redefines the region as a normal region. If mark-unhighlights is nonzero and Epsilon is already highlighting a normal region, this command turns off the highlighting. See set-mark for a command that always begins defining a new region, even when a region has already been highlighted.

mark-paragraph Alt-h Put point and mark around the paragraph.

This command positions mark before the first character in the current paragraph, and positions point after the last character in the paragraph. You can use this command in conjunction with the kill-region command to kill paragraphs and move them around.

For information on Epsilon's notion of a paragraph, see the help entry for the forward-paragraph command.

mark-rectangle Ctrl-x #, Brief: Alt-c Begin marking a rectangular region.

This command begins marking and highlighting a rectangular region of text, setting mark equal to the value of point. A rectangular region consists of all columns between those of point and mark, on all lines in the buffer between point and mark.

If Epsilon is already highlighting a region of another type, this command redefines the region as a rectangular region. If mark-unhighlights is nonzero and Epsilon is already highlighting a rectangular region, this command turns off the highlighting.

mark-whole-buffer Ctrl-x h Highlight the entire buffer.

This command sets point at the start of the current and mark at its end, and turns on highlighting.

maybe-show-matching-delimiter Insert character and show match.

This command first invokes normal-character to insert the key that invoked it, then, if the Matchdelim variable is nonzero, shows the delimiter character matching this one using find-delimiter.

merge-diff Use #ifdef to mark buffer changes.

This command is a variation on the diff command that's useful when comparing files in a C-like language. It marks differences by surrounding them with #ifdef preprocessor lines, first prompting for the #ifdef variable name to use. The resulting buffer receives the mode and settings of the first of the original buffers.

mouse-center M-(Center) Pan or yank, as appropriate.

This command runs either mouse-yank to yank text, mouse-pan to provide panning, or does nothing. See the variable mouse-center-yanks to customize this behavior. By default this command yanks if you press Shift, and pans if not.

mouse-move M-⟨Move⟩ Pop up a scroll bar or menu bar as needed.

Epsilon runs this command when you move the mouse. It pops up a scroll bar or menu bar, or changes the mouse cursor's shape, based on the mouse's current position on the screen.

mouse-pan M-⟨Center⟩ Autoscroll or pan the current buffer.

This command is bound to the middle mouse button on three button (or wheeled) mice. It provides autoscrolling and panning when you click that button.

mouse-select M-⟨Left⟩ Select text, move borders,
or run menu command.

Press and release this mouse button to position point to wherever the mouse cursor indicates, switching windows if needed. Hold down the mouse button and drag to select and highlight text. Double-clicking selects full words. (When a pop-up list of choices appears on the screen, double-clicking on a choice selects it.) Shift-clicking extends the current selection. Holding down the Alt key while selecting produces a rectangle selection.

Drag selected text to move it someplace else. Hold down the Control key to copy the text someplace else.

On scroll bars, this button scrolls the window. You can drag the central scroll box up and down, click on the arrows at the top and bottom of the scroll bar to scroll by lines, or click between the arrows and the box to scroll by pages.

On other window borders and corners, dragging resizes windows. For pop-up windows only, dragging the title bar moves the window.

mouse-to-tag Go to the definition of the indicated function.

If you bind this command to a mouse button, then double-clicking it on the name of a subroutine will jump to that routine's definition using the tags system. Before jumping, it sets a bookmark at the current position like the `set-bookmark` command. Single-clicking positions point and selects text. On window borders, it resizes windows.

For example, these `einit.ecm` commands may be used to bind this command to the right mouse button:

```
~reg-tab "M-<Right>": mouse-to-tag
~reg-tab "M-<DRight>": mouse-to-tag
```

mouse-yank Unix: M-⟨Center⟩ Yank from the clipboard or kill buffer.

This command yanks text from the clipboard or a kill buffer, like the `yank` command, at the mouse's current location.

move-to-window Ctrl-x ⟨Arrows⟩ Move to a different window.

This command changes the current window to the window in the direction of the arrow key from the cursor. For example, typing `Ctrl-x ⟨Right⟩` moves to the window to the right of the cursor; `Ctrl-x ⟨Left⟩` moves to the left. `Ctrl-x ⟨Up⟩` and `Ctrl-x ⟨Down⟩` move up and down, respectively.

`name-kbd-macro` Name the last keyboard macro defined.

Use this command to give a name to a keyboard macro that you defined with `start-kbd-macro` and `end-kbd-macro`. The command prompts you for the name. Thereafter, you may invoke that macro by name using `named-command`, or bind it to a key using `bind-to-key`. Also see `bind-last-macro`, which combines this command with binding.

`named-command` Alt-x Invoke the given command by name.

This command prompts for the name of a command or keyboard macro, with completion, then executes it.

`narrow-to-region` Temporarily restrict editing to between point and mark.

This command temporarily restricts your access to the current buffer. Point can only vary between the values point and mark had when you invoked `narrow-to-region`. The commands that go to the beginning and end of the buffer will instead go to the beginning and end of this region. Searches will only operate within the region. However, the commands that write the buffer to a file will write the entire buffer, not just the constricted region. See also the `widen-buffer` command.

`new-file` Create an empty buffer.

This command creates a new, empty buffer and marks it so that Epsilon will prompt for a file name when you try to save it. You can customize the behavior of the `new-file` command by setting the variables `new-file-mode` and `new-file-ext`.

`next-buffer` F12 Select the next buffer.

This command selects the next buffer and connects it to the current window. You can cycle through all the buffers by repeating this command. To cycle in the other direction, use the `previous-buffer` command.

`next-difference` Vdiff: Alt-(Down) or Alt-] Move to the next change.

Use this command in a buffer created by the `visual-diff` command to move to the next group of changed lines, or the next group of common lines. Added lines are shown in yellow, deleted lines in red, and common lines are colored as in the original buffers.

`next-error` Find a compiler error message, then jump to the offending line.

This command searches in the process buffer for a line containing a compiler error message. Epsilon uses a regular expression search to recognize these messages.

If a window displays the file containing the error, Epsilon switches to that window. Otherwise, it uses the `find-file` command to display the file in the current window. It then goes to the indicated line of the file using the `goto-line` command, then displays the error message in the echo area. A positive numeric argument of n moves to the n th next error message. A negative numeric argument of $-n$ moves to the n th previous error message. A numeric argument of zero repeats the last message.

Users running Cygwin tools may wish to set the `cygwin-filenames` variable to make Epsilon recognize file names in this format.

`next-match` Go to the next matching line.

This command moves to the next match that the last `grep` command found.

If a window displays the file containing the match, Epsilon switches to that window. Otherwise, it uses the `find-file` command to display the file in the current window. It then goes to the matching line. A positive numeric argument of n moves to the n th next match. A negative numeric argument of $-n$ moves to the n th previous match. A numeric argument of zero goes to the same match as last time.

`next-page` Ctrl-v Display the next window full of text.

This command scrolls the current window up so that the last few lines appear at the top of the window. It moves point so that it appears centered vertically in the window.

`next-position` Ctrl-x Ctrl-n Go to the next matching line.

This command moves to the next compiler error message by calling `next-error`, or to the next match found by the `grep` command by calling `next-match`, depending on whether you've run a process or compilation command, or a `grep` command, most recently.

If a window displays the file containing the match, Epsilon switches to that window. Otherwise, it uses the `find-file` command to display the file in the current window. It then goes to the appropriate line of the file. A positive numeric argument of n moves to the n th next match. A negative numeric argument of $-n$ moves to the n th previous match. A numeric argument of zero goes to the same place as last time.

`next-tag` Ctrl-(NumPlus) Go to the next tag with this name.

After you use the `goto-tag` or `pluck-tag` commands to go to a tag that occurs in multiple places, you can use this command to go to the next instance of the tag.

`next-window` Alt-(End), Ctrl-x n Move to the next window.

This command moves to the next window, wrapping around to the first window if invoked from the last window.

You can think of the window order as the position of a window in a list of windows. Initially only one window appears in the list. When you split a window, the two child windows replace it in the list. The top or left window comes before the bottom or right window. When you delete a window, that window leaves the list.

`normal-character` Insert the invoking key into the buffer.

When you type a character bound to the `normal-character` command, Epsilon inserts the character into the buffer, generally before point. See also the `overwrite-mode` command.

Nothing happens if the key that invokes `normal-character` does not represent a valid 8-bit ASCII character.

During auto fill mode, when you type a key bound to this command, the line breaks if appropriate. In particular, if point's column equals the fill column, the command breaks the line. If the value of point's column exceeds the fill column, the command breaks the line at the closest whitespace to the left of the fill column, and uses the `normal-character` command to insert a space. Otherwise, this command just invokes the `normal-character` command to insert the key into the buffer. See the `auto-fill-mode` command.

nsis-mode A mode for editing NSIS files.

This mode provides syntax highlighting suitable for scripts written in the NSIS scripting language.

oem-to-ansi Convert buffer's DOS character set to Windows.

Windows programs typically use a different character set than do DOS programs. The DOS character set is known as the DOS/OEM character set, and includes various line drawing characters and miscellaneous characters not in the Windows/ANSI set. The Windows/ANSI character set includes many accented characters not in the DOS/OEM character set. Epsilon for Windows uses the Windows/ANSI character set (with most fonts).

The oem-to-ansi command converts the current buffer from the DOS/OEM character set to the Windows/ANSI character set. If any character in the buffer doesn't have a unique translation, the command warns first, and moves to the first character without a unique translation.

This command ignores any narrowing established by the narrow-to-region command. It's only available in Epsilon for Windows.

one-window Ctrl-x 1 Display only one window.

The current window becomes the only window displayed. The buffers associated with other windows, if any, remain unaffected. See also the zoom-window command.

open-line Ctrl-o Open up some vertical space.

This command inserts a newline after point, instead of before point as the other inserting commands do. Use this command to open up some vertical space in the file.

overwrite-mode <Ins> Enter/Exit overwrite mode.

This command changes the behavior of the normal-character command, causing it to insert characters into the buffer destructively, replacing the character after point. However, Epsilon will never overwrite a newline character, or overwrite another character with a newline character. This ensures proper behavior with respect to the ends of lines.

Without a numeric argument, the command toggles the state of overwrite mode. With a numeric argument of zero, the command disables overwrite mode. With a nonzero numeric argument, it turns overwrite mode on.

page-left Alt-⟨PageUp⟩ Show more text to the left.

This command moves left on the current line by half the window's width. In horizontal scrolling mode, it then horizontally centers the window on that column (or, if possible, positions the window so that the start of the line is also visible).

page-right Alt-⟨PageDown⟩ Show more text to the right.

This command moves right on the current line by half the window's width. In horizontal scrolling mode, it then horizontally centers the window on that column (or, if possible, positions the window so that the start of the line is also visible).

`prefix-fill-paragraph` `Alt-Shift-Q` Fill a paragraph preserving indentation and similar.

This command examines the current line to determine its “prefix”, a run of non-alphanumeric characters at the start of the line. Then it determines the boundaries of the current paragraph among those adjacent lines that share this prefix. Then it fills that paragraph, ensuring the prefix remains at the start of each line. With a region highlighted, it fills all paragraphs within the region in the same manner.

With a numeric argument, it fills paragraphs using the current column as the right margin, instead of the `margin-right` variable.

The `mail-fill-paragraph` command is similar, but specialized for email quoting rules.

`previous-buffer` `F11` Select the previous buffer.

This command selects the previous buffer and connects it to the current window. You can cycle through all the buffers by repeating this command. To cycle in the other direction, use the `next-buffer` command.

`previous-difference` `Vdiff: Alt-(Up)` or `Alt-[` Move to the previous change.

Use this command in a buffer created by the `visual-diff` command to move to the start of the previous group of changed lines, or the previous group of common lines. Added lines are shown in yellow, deleted lines in red, and common lines are colored as in the original buffers.

`previous-error` Find a compiler error message, then jump to the offending line.

This command works like `next-error`, except that it searches backwards instead of forwards.

`previous-match` Go to the previous matching line.

This command moves to the previous match from the last `grep` command.

If a window displays the file containing the match, `Epsilon` switches to that window. Otherwise, it uses the `find-file` command to display the file in the current window. It then goes to the matching line. A positive numeric argument of n moves to the n th previous match. A negative numeric argument of $-n$ moves to the n th next match. A numeric argument of zero goes to the same match as last time.

`previous-page` `Alt-v` Display the previous window full of text.

This command scrolls the contents of the current window down so that the first few lines appear at the bottom of the window. It moves point so that it appears centered vertically in the window.

`previous-position` `Ctrl-x Ctrl-p` Go to the previous matching line.

This command moves to the previous compiler error message by calling `previous-error`, or to the previous match found by the `grep` command by calling `previous-match`, depending on whether you’ve run a process or compilation command, or a `grep` command, most recently.

If a window displays the file containing the match, `Epsilon` switches to that window. Otherwise, it uses the `find-file` command to display the file in the current window. It then goes to the appropriate line of the file. A positive numeric argument of n moves to the n th previous match. A negative numeric argument of $-n$ moves to the n th next match. A numeric argument of zero goes to the same place as last time.

previous-tag Ctrl-⟨NumMinus⟩ Go to the previous tag with this name.

After you use the goto-tag or pluck-tag commands to go to a tag that occurs in multiple places, you can use this command to go to the previous instance of the tag.

previous-window Alt-⟨Home⟩, Ctrl-x p Move to the previous window.

This command moves to the previous window, wrapping around to the last window if invoked from the first window.

You can think of the window order as the position of a window in a list of windows. Initially only one window appears in the list. When you split a window, the two child windows replace it in the list. The top or left window comes before the bottom or right window. When you delete a window, that window leaves the list.

print-buffer Alt-F9 Print the current buffer.

This command prints the current buffer. Under Windows, it displays the standard Windows printing dialog. You can choose to print the current selection, the entire buffer, or just certain pages.

Under other environments, this command prints the current highlighted region. If no region in the buffer is highlighted, the command prints the entire buffer.

It prompts for the device name of your printer, storing your response in the variable `print-destination` (or, under Unix, `print-destination-unix`), and then writes a copy of the buffer to that device.

If the printer name begins with the `!` character, Epsilon interprets the remainder of the name as a command line to execute in order to print a file. Epsilon substitutes the file to be printed for any `%f` sequence in the command line. For example, if your system requires you to type “netprint filename” to print a file, enter `!netprint %f` as the device name and Epsilon will run that command, passing it the file name of the temporary file it generates holding the text to print. The device name can include any of the file name template sequences, such as `%p` for the path to the file to print. It can also include `%t`, which substitutes the name of the buffer or file being printed (which will be different from the name of temporary file with the text to print). It may be used as a title or heading.

If the variable `print-tabs` is zero, Epsilon will make a copy of the text to print and convert any tabs into spaces before sending it to the printer.

print-buffer-no-prompt Print the current buffer without prompting.

This command prints the current buffer, exactly like `print-buffer`, but doesn’t prompt. It uses default settings.

print-region Shift-F9 Print the current region.

Under Windows, this command displays the standard Windows printing dialog. You can choose to print the current selection, the entire buffer, or just certain pages.

Under other environments, this command always prints the current region.

See the `print-buffer` command for details on printing.

`print-setup` Display the Print Setup dialog.

Under Windows, this command displays the standard Print Setup dialog. You can choose a printer and select other options. In other environments, this command does nothing.

`process-backward-kill-word` Process mode: `Ctrl-Alt-h` Kill the word before point.

The command moves point as in `backward-word`, killing the characters it passes over. But it stops before deleting any part of the prompt, treating that as a word boundary.

`process-complete` Process mode: `<Tab>` Finish typing a file name.

In a process buffer, `<Tab>` performs completion on file names. If no more completion is possible, it displays all the matches in the echo area, if they fit. If not, press `<Tab>` again to see them listed in the buffer. See the `process-completion-style` variable to customize how process completion works. The variables `process-completion-dircmds` and `process-completion-windows-programs` provide additional customization options.

The command uses different rules for the first word on the command line, searching for a command along the `PATH` in a manner appropriate to the operating system. (It won't know about any commands that may be built into the current shell command processor, though.)

At buffer positions before the prompt, this command indents.

`process-enter` Process mode: `<Enter>` Send a line to the concurrent process.

Pressing the `<Enter>` key in process mode moves the error spot backwards to point, so that Epsilon searches for error messages from this new location. If the `process-enter-whole-line` variable is nonzero, Epsilon moves to the end of the current line before sending it to the process, but only when in a line that has not yet been sent to the process. If the `process-enter-whole-line` variable is two, Epsilon copies the current line to the end of the buffer, making it easier to repeat a command.

`process-mode` Interact with a concurrent process.

Epsilon puts its process buffer in this mode. Pressing the `<Enter>` key in process mode moves the error spot backwards to point, so that Epsilon searches for error messages from this new location. Process mode also includes commands for completing on file names and command names (using the `process-complete` command on `<Tab>`) and retrieving previous command lines (using the `process-previous-cmd` and `process-next-cmd` commands on `Alt-p` and `Alt-n`).

`process-next-cmd` Process: `Alt-n` Retrieve the next command from the history list.

Epsilon's concurrent process buffer maintains a command history. This command retrieves the next command from the history. Use it following a `process-previous-cmd` command. With a numeric prefix argument, the command shows a menu of previous commands and you can select one to repeat.

process-previous-cmd	Process: Alt-p	Retrieve the previous command from the history list.
----------------------	----------------	--

Epsilon's concurrent process buffer maintains a command history. This command retrieves the previous command from the history. Also see `process-next-cmd` command. With a numeric prefix argument, the command shows a menu of previous commands and you can select one to repeat.

process-yank	Process mode: Ctrl-y	Insert the contents of a kill buffer.
--------------	----------------------	---------------------------------------

This command behaves just like the `yank` command, but if more than one line would be yanked (and then immediately executed by the running shell command processor), it first prompts for confirmation. When a keyboard macro is running or being defined, this prompting is disabled.

profile	Collect timing information on EEL commands.
---------	---

This command starts a recursive edit and begins collecting timing data. Many times per second, Epsilon makes a note of the currently executing EEL source line. When you exit with `exit-level`, it fills a buffer named "profile" with this timing data. Epsilon doesn't collect any profiling information on commands or subroutines that you compile with the `-s` option.

pull-word	F3, Ctrl-⟨Up⟩	Complete this word by scanning the buffer.
-----------	---------------	--

This command scans the buffer before point, and copies the previous word to the location at point. If you type the key again, it pulls in the word before that, etc. Whenever Epsilon pulls in a word, it replaces any previously pulled-in word. If you like the word that has been pulled in, you do not need to do anything special to accept it—Epsilon resumes normal editing when you type any key except for the few special keys reserved by this command. Type Ctrl-g to erase the pulled-in word and abort this command.

If a portion of a word immediately precedes point, that subword becomes a filter for pulled-in words. For example, suppose you start to type a word that begins `WM`, then you notice that the word `WM_QUERYENDSESSION` appears a few lines above. Just type Ctrl-⟨Up⟩ and Epsilon fills in the rest of this word.

pull-word-fwd	Ctrl-⟨Down⟩	Complete this word by scanning the buffer.
---------------	-------------	--

This command scans the buffer after point, and copies the next word to the location at point. If you type the key again, it pulls in the word after that, etc. Whenever Epsilon pulls in a word, it replaces any previously pulled-in word. If you like the word that has been pulled in, you do not need to do anything special to accept it—Epsilon resumes normal editing when you type any key except for the few special keys reserved by this command. Type Ctrl-g to erase the pulled-in word and abort this command.

If a portion of a word immediately precedes point, that subword becomes a filter for pulled-in words. For example, suppose you start to type a word that begins `WM`, then you notice that the word `WM_QUERYENDSESSION` appears a few lines below. Just type Ctrl-⟨Up⟩ and Epsilon fills in the rest of this word.

python-indenter	⟨Tab⟩ in Python mode	Indent this line for Python.
-----------------	----------------------	------------------------------

In a line's indentation, reindent the line correctly for Python code. Inside the text of a line, or when repeated, insert a tab.

If a region is highlighted, Epsilon indents all lines in the region by one tab stop. With a numeric prefix argument, Epsilon indents by that amount.

python-mode

Set up for editing programs in the Python language.

This command puts the current buffer in a mode suitable for editing programs in the Python language. Syntax highlighting, indenting, tagging, comment filling, delimiter highlighting and commenting commands are all provided.

push

Ctrl-x Ctrl-e

Invoke an inferior command processor.

This command invokes a command processor, or shell. While in the command processor, you can do whatever you normally do outside of Epsilon. If you prefix a numeric argument, the command prompts you for a line to pass to the command processor. After you have passed a command to the command processor, that command becomes the default command until you type in another one.

Unlike start-process, this command will work with all but the most extremely misbehaved programs.

query-replace

Alt-%

Interactively replace strings.

This command behaves like `replace-string`. Instead of replacing everything automatically, it positions point after each occurrence of the old string, and you may select whether or not to replace it. With a numeric argument, the command will match only complete words.

y or **<Space>** replaces and goes to the next match.

n or **<Backspace>** doesn't replace, but goes to the next match.

<Esc> exits immediately.

. replaces and then exits.

^ backs up to the previous match.

! replaces all remaining occurrences.

, replaces the current match but doesn't go to the next match.

Ctrl-r enters a recursive edit, allowing you to modify the buffer arbitrarily. When you exit the recursive edit with `exit-level`, the query-replace continues.

Ctrl-g exits and returns point to its original location.

Ctrl-w toggles the state of word mode.

Ctrl-t toggles the state of regular expression mode (see the description of `regex-replace`).

Ctrl-c toggles the state of case-folding.

Any other key causes query-replace to exit and any command bound to that key to execute.

See the `replace-in-region` variable to restrict matches to a highlighted region.

quick-dired-command

Alt-o

Perform operations on the current file.

This command provides a convenient way to perform various operations on the file associated with the current buffer. It prompts for another key, with choices as listed below. Many of them are similar to the corresponding commands in a dired buffer.

d deletes the file associated with the current buffer, after prompting for confirmation.

`rebuild-menu` Put modified bindings into menu.

This command makes Epsilon reconstruct its menus, adding current key bindings.

`record-kbd-macro` Brief: F7 Start or stop recording a macro.

This command begins recording a keyboard macro. Keys you press execute normally, but also become part of an accumulating keyboard macro. Run this command again to finish defining the macro.

`redisplay` Rewrite the entire screen.

Normally Epsilon does not write to the screen during the execution of a keyboard macro. This command forces a complete rewrite of the screen. Use it if you need to create a keyboard macro that updates the screen in the middle of execution.

`redo` F10 Redo the last buffer change or movement.

This command reverses the effect of the last `undo` command. If repeated, it restores earlier changes. You may remove the changes again with `undo`.

`redo-by-commands` Redo, grouping by commands.

This command operates like `redo`, except that it groups editing operations based on commands, similar to Brief.

`redo-changes` Ctrl-F10 Redo, grouping movements.

This command operates like `redo`, except that it will automatically redo any run of changes to the buffer that involve only movements of point as a unit, and stop just before a change of actual buffer contents. When you invoke `redo-changes`, it performs a `redo`, and if the redone operation simply moved point, then it continues to redo changes until it encounters a non-movement operation.

`redo-movements` Ctrl-F12 Return to the site of previous edits.

After making changes at various spots in a buffer, use the `undo-movements` command and this command to move to the location of previous edits. Epsilon uses undo information to record where in a buffer editing has occurred. This command moves forward through the list of locations, toward more recent locations.

`refresh-files` Check to see which files have changed on disk.

This command makes Epsilon check each buffer to see if its associated file has been modified on disk. It rereads the modified file automatically, or asks permission to do so, just as if you had switched to every buffer one by one.

regex-replace

Alt-*

Substitute for replace expressions.

This command functions like query-replace, but starts in regular expression mode.

pat1|pat2 matches either *pat1* or *pat2*.
*pat** matches zero or more matches of *pat*.
pat+ matches one or more matches of *pat*.
pat? matches zero or one matches of *pat*.
[abx] matches any of the characters a, b, or x.
[^abx] matches any but a, b, or x.
[a-z3] matches a, b, c, ... z, or 3.
. matches any character except newline.
() group patterns for +, *, ?, and |.
^ only matches at the beginning of a line.
\$ only matches at the end of a line.
<#50> means the character with ASCII code 50.
% removes the special meaning from the following character, so that %\$ matches only \$.
! marks the end of the match. The command does not change any characters that match the pattern after the exclamation point.

In the replacement text, #1 means substitute the part of the text that matched the first parenthesized pattern piece. For example, using regex-replace to replace “([A-Z] [a-z]+) ([.!?])” with “#2 ends #1” changes the text “Howard!” to “! ends Howard”. #0 means to substitute the whole match. #U forces any following replacement text to uppercase, #L and #C to lowercase or capitalized. #E ends such case modifications; the remaining replacement text will be substituted as-is. #S substitutes the next alternative, when the search pattern consists of simple alternative bits of fixed text separated by |’s. Characters may be included by name in replacement text using the syntax #<Newline>, which substitutes a <Newline> character.

regex-search

Ctrl-Alt-s

Search for a string after point.

The command prompts for a regular expression, then positions point after the next match of that pattern. If no such match exists, a message appears in the echo area.

pat1|pat2 matches either *pat1* or *pat2*.
*pat** matches zero or more matches of *pat*.
pat+ matches one or more matches of *pat*.
pat? matches zero or one matches of *pat*.
[abx] matches any of the characters a, b, or x.
[^abx] matches any but a, b, or x.
[a-z3] matches a, b, c, ... z, or 3.
. matches any character except newline.
() group patterns for +, *, ?, and |.
^ only matches at the beginning of a line.
\$ only matches at the end of a line.
<#50> means the character with ASCII code 50.
% removes the special meaning from the following character, so that %\$ matches only \$.
! marks the end of the match.

`release-notes` Display the release notes.

Epsilon searches for the file `readme.txt` and loads it.

`rename-buffer` Change the name of the current buffer.

Epsilon prompts for a buffer name, then renames the current buffer.

`replace-again` Brief: Shift-F6 Do the last replacement again.

This command repeats the last replace command you did, using the same text to search for, and the same replacement text.

`replace-string` Alt-& Replace one string with another.

The command asks you for the old and new strings. From point to the end of the buffer, it replaces occurrences of the old string with the new string. If you prefix a numeric argument, it will only replace matches that consist of complete words. See also `query-replace`.

`reset-mode` Pick the appropriate mode for this buffer.

When you first load a file, Epsilon auto-detects the correct mode for it, by examining the file's extension and sometimes the contents of the file. This command makes Epsilon repeat that process, setting the buffer to a different mode if appropriate. It can be handy after you've temporarily switched to a different mode for any reason, or after you've started creating a new file with no extension and have now typed the first few lines, enough for Epsilon to auto-detect the proper mode.

For instance, if you're creating a new file with no extension, there might not be enough information for Epsilon to choose the right mode at the start. Once you've typed the usual first line of a Perl, PostScript, shell script, or similar file, then Epsilon should have enough information to pick the right mode.

`resume-client` Ctrl-c # Tell a waiting client you've finished editing.

You can set up Epsilon so an external program can run it as its editor. In Unix-style environments this is typically done by setting the `EDITOR` environment variable. The external program will invoke the editor program, and then wait for it to exit before continuing with its work.

You may have an existing session of Epsilon running, and want all editing requests from other programs to be routed to the existing session. You can set that up with Epsilon by setting `EDITOR` to `epsilon -wait`. The external program will run a second copy of Epsilon (the client), which will pass the name of the file to be edited to the existing Epsilon session (the server), and wait for the server before continuing.

When you've finished editing the passed file, save it, and then use the `resume-client` command to notify the client instance of Epsilon that the editing job is done, and it should exit.

In Windows, there's no convention similar to `EDITOR`, but individual programs may have an option to set the name of the editor program to be used. Note that by default, the Windows command processor never waits for a GUI-style program to exit before prompting again; putting `"start /wait"` before the program name tells it to wait. So `"start /wait epsilon -wait"` may be needed.

retag-files Tag all files again.

This command retags all files in mentioned in the current tag file.

reverse-incremental-search Ctrl-r Incremental search backwards.

This command starts incremental-search in reverse.

reverse-regex-search Ctrl-Alt-r Search for a string before point.

This command prompts for a regular expression, then positions point before the first match of that string before point. If no such match exists, a message appears in the echo area.

reverse-replace Interactively replace strings, moving backward.

This command behaves like query-replace, but searches backward through the buffer for text to replace, instead of forward. It positions point before each occurrence of the old string, and you may select whether or not to replace it. With a numeric argument, the command will match only complete words.

reverse-search-again Search backward for the same search string.

reverse-sort-buffer Reverse sort the current buffer.

This command asks for the name of a buffer and fills it with a copy of the current buffer reverse sorted by lines. If you specify a numeric argument of n , the command will ignore the first n columns on each line when comparing lines. You can press Alt-g at the prompt to insert the current buffer's name, to make it easier to sort the current buffer in place.

reverse-sort-region Reverse sort part of the buffer in place.

This command reverse sorts in place the lines of the current buffer appearing between point and mark. If you specify a numeric argument of n , the command will ignore the first n columns on each line when comparing lines.

reverse-string-search Reverse search in non-incremental mode.

This command starts a reverse search in non-incremental mode. It functions like starting a reverse-incremental-search, then disabling incremental searching with Ctrl-o.

revert-file Read the current file into this buffer again.

Epsilon replaces the contents of the current buffer with the contents of the current file on disk. If the current buffer has unsaved changes, Epsilon asks if you want to discard the changes by reading the file.

run-with-argument Prompt for a numeric argument, then get key and run.

This command prompts for a numeric argument, then passes it to the next command you invoke. It's similar to the argument command, but it lets you enter an argument saved in a kill buffer (say, inside a keyboard macro), or enter the argument in a different base (e.g. 0x100), or using any other syntax Epsilon recognizes when prompting for a number.

Then it asks if you want to restrict searching to particular man page sections, such as 1 for commands or 3 for subroutines. Finally, it asks if you want to restrict the search to man page entries matching a certain file pattern, such as `*file*` to search only pages whose names contain “file”.

For speed reasons, it searches each man page without processing it through the `man` command, searching the man page in source format. By default, it shows only the first match in each page; set the `search-man-pages-shows-all` variable to see all matches. The result appears in the `grep` buffer; when you view a match from there, Epsilon will then use the `man` command to display its processed form.

search-region Search for text only in the current region.

This command begins a search that is restricted to the region between point and mark. It works with rectangular and line regions too. Epsilon will highlight the region until the command ends.

This command is a variation on the `incremental-search` command, and all its special keys operate as usual. In particular, you can use `Ctrl-r` and `Ctrl-s` to change the direction of the search.

You can set the `search-in-region` variable to make this type of restricted search happen whenever you highlight a region and then search using the usual search commands.

select-and-grep Search multiple files with fewer restrictions.

This command prompts to select from a predefined `grep` rule set, then runs `grep` using that rule set. A rule set provides an alternate set of values for one or more of the variables

`use-grep-ignore-file-variables`, `grep-ignore-file-extensions`,
`grep-ignore-file-pattern`, `grep-ignore-file-basename`,
`grep-ignore-file-types`, or `grep-ignore-file-max-size`.

By default, the only rule set is named “all”, and it provides a setting of zero for `use-grep-ignore-file-variables`, meaning it makes Epsilon search all files, not skipping any based on their name or other characteristics.

You can define a new rule set by simply defining a variable whose name is of the form *varname-rulesetname*. (For example, the default rule set was created by defining a variable named `use-grep-ignore-file-variables-all`.) The rule set variable should be set to the value for that variable when the rule is applied. For example, you could make a rule that skips only `.bak` files by defining a character array variable named `grep-ignore-file-extensions-nobakfiles` and setting it to “|.bak|”. The rule set name “nobakfiles” will then automatically appear in this command’s rule set list.

select-buffer `Ctrl-x b` Display a buffer in the current window.

This command prompts for a buffer name. If a buffer with that name exists, the command connects it to the current window. Otherwise, it creates a buffer with the indicated name, and connects that to the current window. If you just press `(Enter)` to the prompt for a buffer, it defaults to the last buffer associated with the current window. So, repeated `Ctrl-x b`’s will generally switch back and forth between two buffers.

select-browse-file Pick a `.bsc` file for the source code browser interface.

This command selects the Microsoft browser database file to be used by the `browse-symbol` command. It sets the `browser-file` variable to the chosen `.bsc` file name. If you don’t use this command to select a `.bsc` file, but you have selected a `.bsc` file for tagging using the `select-tag-file` command, Epsilon will use it for browsing too.

The first group specifies which graphic characters Epsilon should use to draw window borders. It defines all the line-drawing characters needed for drawing four different styles of borders, and all possible intersections of these. Epsilon for Windows does not use this group.

The next group specifies which characters Epsilon uses to display various special characters like `<Tab>` or Control-E. For example, Epsilon usually displays a control character with the `^` symbol. Set the appropriate character in this group to make Epsilon use a different character. You can also make Epsilon display a special character at the end of each line, or change the continuation character. (Commands like `change-show-spaces` also set some of these positions.)

The last group defines the characters Epsilon uses to display window scroll bars. Epsilon replaces the window's selected border characters with characters from this group. Epsilon for Windows does not use this group.

`set-display-look`

Make the screen look like another editor.

This command makes Epsilon's window decoration and screen appearance resemble that of some other editor. It displays a menu of choices. You can select Epsilon's original look, Brief's look, the look of the DOS Edit program (which is the same as the QBasic program), or the look of Borland's IDE.

`set-encoding`

Make Epsilon use a different encoding.

When Epsilon reads or writes a file that uses Unicode characters, it uses an encoding such as UTF-8 or UTF-16 to represent the file on disk. This command sets the encoding Epsilon will use from now on for the current buffer. Press `“?”` to see a list of known encodings.

Some encodings are special. The `“raw”` encoding makes Epsilon read a file without doing any conversion. The `“auto-detect”` encoding makes Epsilon determine the encoding as it reads the file, by looking for a signature (byte order mark). Some encodings like UTF-8 have `“-No-Sig”` variants that omit the signature when writing.

To read a file into a new buffer using specific encoding, use a numeric prefix argument with a file-reading command like `find-file`.

`set-file-name`

Brief: Alt-o

Change the file name associated with this buffer.

This command prompts for a new file name for the current buffer, and changes the file name associated with the buffer. The next time you save the file, Epsilon will save it under the new name.

`set-fill-column`

Ctrl-x f

Set the column at which filling occurs.

If you provide a numeric argument, the command sets the fill column for the current buffer to that value. Otherwise, the command prompts you for a new fill column, with the point's column offered as a default. The fill column controls what auto fill mode and the filling commands consider the right margin.

To set the default value for new buffers you create, use the `set-variable` command on F8 to set the default value of the `margin-right` variable. (Or for C mode buffers, set the `c-fill-column` variable.)

set-font Select a different font.

This command changes the font Epsilon uses, by displaying a font dialog box and letting you pick a new font. It's available under Windows and X11.

To select a Unicode font on X11 and Mac systems, first select iso10646-1 in the list of character sets on the Filter pane of the font selection dialog.

set-line-translate Specify Epsilon's line translation scheme.

The operating system uses the sequence of characters Return Newline to indicate the end of a line. Epsilon normally changes this sequence to a single Newline when it reads in a file (by removing all the Return characters). When it writes a file, it adds a Return before each Newline character.

Epsilon automatically selects one of several other translation types when appropriate, based on the contents of the file you edit (regular text, binary, Unix, or Macintosh). You can explicitly override this if Epsilon guesses wrong by providing a numeric argument to a file reading command like `find-file`. Epsilon will then prompt for which translation scheme to use.

This command sets the desired translation method for the current buffer. It prompts for the desired type of translation, and makes future file reads and writes in this buffer use that translation. Epsilon will display "Binary", "Unix", "DOS", or "Mac" in the mode line to indicate any special translation in effect.

set-mark Ctrl-@ Set the mark to the current position.

Commands that operate on a region of the buffer use the mark and point to delimit the region. This command sets the mark to the current value of point.

set-named-bookmark Ctrl-x / Name the current editing position.

This command prompts you for a letter, then associates that letter with a bookmark at the current location. Subsequently, you can return to that location with the `jump-to-named-bookmark` command. If you provide a digit instead of a letter, Epsilon sets the appropriate temporary bookmark (0 refers to the last one, 1 to the one before that, and so on). You can press '?' to get a list of the currently defined bookmarks, along with the text that contains the bookmarks. To select one, simply move to the desired bookmark and press `(Enter)`. See also `set-bookmark` and `jump-to-last-bookmark`.

set-printer-font Select the font to use when printing.

Use this command to select the font Epsilon uses when printing. It sets the variable `font-printer`.

set-show-graphic Show codes or glyphs for some characters.

By default, Epsilon displays most control characters (codes below 32) by prefixing to them a caret, e.g., Control C appears as "`^C`". It displays other characters, including national characters, with their graphic symbol (glyph). Epsilon has four different modes for displaying all these characters.

In mode 0, Epsilon displays Meta characters (characters with codes from 160 to 255, meaning ASCII characters but with the 8th bit on) by prefixing to them a "`M-`", e.g., Meta C appears as

“M-C”. Epsilon display Control-meta characters (128 to 159) by prefixing to them “M-^”, e.g., “M-^C”. Epsilon displays most control characters (codes from 0 to 31) by prefixing to them a caret, e.g., Control C appears as “^C”. Other characters are shown with their usual glyph.

In mode 1, Epsilon uses glyphs to display all the above characters (except for the few that have a special meaning, like <Tab> or <Newline>).

In mode 2, hex mode, Epsilon displays all characters outside the range 32 to 127 (including control and meta characters) by their hexadecimal ASCII values, with an “x” before them to indicate hex.

In mode 3, which is the default, Epsilon displays control characters as “^C”, and shows glyphs for all other characters, as described above.

If you provide no numeric argument, this command cycles to the next mode in the above list. A numeric argument of 0, 1, 2, or 3 selects the corresponding mode.

set-tab-size Set how many columns are between tab settings.

This command sets the number of spaces between tab stops for the current buffer. If given a numeric argument, Epsilon sets the tab size to that number. Otherwise the command prompts for the tab size. By default, Epsilon puts tab settings every 8 columns. Some language modes like C mode default to a different setting; see `c-tab-override` and similarly-named variables. This command will offer to set one of those too if appropriate.

set-variable F8 Set any EEL variable.

This command prompts for the name of a variable and a new value for that variable. This command cannot set variables with complicated types involving structures or pointers. After setting the variable, Epsilon shows the new value using `show-variable`.

If you specify a buffer-specific or window-specific variable, Epsilon uses the numeric argument to determine whether to set the value for the current buffer or window (zero numeric argument), the default value (1), or both (2). The value 3 also sets all it for all non-system buffers and windows, and 4 includes system ones too. If you provide no numeric argument, Epsilon asks.

As with most prompts for a numeric value, you can enter the value in hexadecimal, octal or binary as 0x2A, 0o52, or 0b101010, respectively, or write an integer expression using the usual arithmetic operators +, -, *, / and so forth.

set-want-backup-file Brief: Ctrl-w Turn backup files on or off in this buffer.

This command toggles whether or not Epsilon makes a backup file each time you save the current buffer.

shell-command Alt-! Run a program and collect its output.

This command prompts for the name of an external program to run. Any output it produces goes to a buffer named “filter-output”. With a negative prefix argument, it prompts for the name of a destination buffer instead. The program runs asynchronously.

With a non-negative numeric prefix argument, the program runs synchronously, and the output is inserted into the current buffer.

shell-mode Set up for editing shell scripts.

This command puts the current buffer in a mode suitable for editing Unix shell scripts and similar files.

show-bindings F5, F1 B Find a key bound to a command.

The command prompts for a command name, then displays a message telling which keys, if any, run that command.

show-connections Ctrl-Alt-c Show all Internet connection buffers.

This command lists all active Telnet, FTP, and similar Internet activities and buffers. You can select a buffer and press `<Enter>` to switch to it, or press `<Escape>` to remain in the current buffer.

show-last-keys F1 L Display recently typed keys.

This command pops up a window that displays the last 60 keystrokes you typed.

show-matching-delimiter Insert character and show match.

This command first invokes `normal-character` to insert the key that invoked it, then shows the delimiter character matching this one using `find-delimiter`. Some people like to bind this command to keys such as `"")` or `"}"`.

show-menu Alt-F2 Display a menu of commands.

This command displays a menu of commands and lets you choose one. Use the arrow keys to navigate through the menu. Letter keys move to the next command in the current column beginning with that letter. Press `<Enter>` to execute the highlighted command, or click on a command with the mouse. Press `Ctrl-g` or `<Esc>` to exit from the menu.

show-point Ctrl-x = Show information about point.

This command displays the column number, value of point, and size of the buffer, as well as the ASCII, decimal, and hex codes of the character after point (plus its name if it's outside the ASCII range). If the character is a Unicode surrogate, it also displays the code and name of the complete Unicode character it helps to form, if any.

The file may occupy more space on disk than the buffer size indicates, due to the line translation scheme that Epsilon uses when reading and writing files, or Unicode encoding. Use the `count-lines` command, bound to `Ctrl-x L`, to get the number of characters the buffer would occupy on disk after line translation. See the `mode-format` variable to display the current column number or the value of point continuously.

show-standard-bitmaps Display available icons for the tool bar.

You can use this function to see some of the icons that may appear on Epsilon's tool bar (Windows GUI version only). It's useful when modifying the contents of the tool bar.

sort-tags Sort the list of tags manually.

By default, Epsilon sorts the tag list whenever it needs to display a list of tag names for you to choose from. Instead, you can set the `want-sorted-tags` variable to 0, and sort the tags manually, whenever you want, using this command.

spell-buffer-or-region Correct misspelled words.

This command moves to each misspelled word in the current buffer, one by one, and prompts, asking if you want to replace with a correction or ignore it. It calls the `spell-correct` command to do this.

With a highlighted region, it operates on only the words within the region.

spell-configure Set up dictionary for spell checking.

This command configures the dictionary file Epsilon uses for spell checking, installing it if necessary. The first time you run it, it will download (if necessary) and unpack a set of dictionary files. See <http://www.lugaru.com/spell.html> if you need to download these manually.

Then it will ask for your preferences on dictionary size, variety of English (American, British, or Canadian spellings), and so forth. (Bigger dictionaries recognize rarer words, but won't catch misspellings that result in rare words.) Run it again to change these preferences.

You can choose to include abbreviations and acronyms like "TV", "VIP", or "rpm", contractions like "doesn't", or "I'm", or words that must start with an uppercase letter, like "Doctor", "December", or "Europe". The proper names category includes less common names like "Brian" or "Altman" that are not included in the uppercase category. (Some words appear in more than one category.)

Variants permit words like "fiascos" in addition to "fiascoes" or "Hanukah" in addition to "Hanukkah".

It also asks if you want to run an external helper program to generate better suggestions. If you have one of the supported programs installed, letting Epsilon use it will produce faster, more appropriate suggestions. Epsilon can use the free `aspell` and `ispell` programs often installed on Unix systems (and available on Windows systems through a Cygwin package from <http://www.cygwin.com>), the commercial `MicroSpell` spell checker for Windows from <http://www.microspell.com> (version 9.1f or later), or the Mac's spelling engine.

This option sets the `spell-helper-program` variable. If your external speller program is not on your `PATH`, you'll have to modify this variable to include its full pathname.

spell-correct Ctrl-c Ctrl-o Replace or ignore a misspelled word.

You can use this command to replace a misspelled word from a list of suggestions, or enter it into an "ignore list" so Epsilon will no longer complain about it.

Your global ignore list is a file named `ignore.lst` in your customization directory. Epsilon also checks directory-specific and file-specific ignore lists. When checking a file named `file.txt`, for example, Epsilon looks for an ignore list file named `.file.txt.espell` in that same directory, and a directory-specific ignore list file in that directory named `.directory.espell`.

All these files contain one word per line. You can edit them to remove words you no longer wish the speller to ignore. Epsilon automatically sorts ignore list files when it uses them.

Epsilon can generate suggestions in one of several ways. Its default method can be replaced by a faster but less accurate method by setting a bit in the appropriate mode-specific variable, such

as `c-spell-options` for C mode or `default-spell-options` for modes without their own variable. Or by setting the `spell-helper-program` variable, using the `spell-configure` command, Epsilon can use an external program to provide suggestions.

`spell-grep` List the misspelled words.

This command makes a list of all lines in the current buffer with a misspelled word, putting its results in the `grep` buffer. You can then select particular lines to correct using the commands of `grep-mode`.

With a highlighted region, it operates on only the words within the region.

`spell-mode` Toggle highlighting of misspelled words.

This command toggles the Spell minor mode. With a nonzero numeric argument, it turns Spell mode on; with a zero argument it turns Spell mode off. In Spell mode, Epsilon highlights misspelled words. The change affects all buffers in the same major mode (and sometimes those in other modes, whenever both modes lack a dedicated `spell-options` variable and use the `default-spell-options` variable for their settings). See the `buffer-spell-mode` command to toggle Spell minor mode for a single buffer, not a group.

Use the `spell-correct` command to ignore certain words or show suggestions for a misspelled word.

`split-window` Ctrl-x 2 Split the current window in two.

This command splits the current window into two windows, one on top of the other, occupying the same total space. Nothing happens if either resulting window would have fewer than two lines of height (counting the mode line).

`split-window-vertically` Ctrl-x 5 Split the current window in two.

This command splits the current window into two windows, one beside the other, occupying the same total space. Nothing happens if either resulting window would have fewer than one character of width.

`ssh` Connect to another computer securely.

This command prompts for the name of a computer, then connects to it by running an `ssh` (secure shell) program on your system. You can use a name such as `bob@example.com` to connect using a specific user name, not the default one determined by the `ssh` program.

`ssh-mode` A mode for secure shell connections.

The `ssh` command creates buffers in this mode, used for connections to a remote host using an `ssh` program. It provides coloring of some ANSI escape sequences, command history using Alt-p and Alt-n, and parsing of password prompts. The variables `ssh-interpret-output` and `recognize-password-prompt` control some of these features. In `ssh` mode, Ctrl-c Ctrl-c sends an interrupt signal to the remote machine.

`standard-toolbar` Display Epsilon's normal tool bar.

Epsilon calls this function to display its tool bar (Windows GUI version only). By redefining the function, you can change what appears on the tool bar.

`start-kbd-macro` `Ctrl-x (` Start defining a keyboard macro.

After you invoke this command, everything you type executes normally, but it also becomes part of an accumulating keyboard macro. The macro definition ends when you invoke the `end-kbd-macro` command.

`start-process` `Ctrl-x Ctrl-m` Invoke a concurrent command processor.

You can create a concurrent subprocess with Epsilon. The `start-process` command shows the “Process” buffer in the current window, and starts a command processor running in it. Epsilon will capture the output of commands that you run in the window, and insert that output into the process buffer. When the process reads input from its standard input, Epsilon will give it the characters that you insert at the end of the buffer. You can move to other windows or buffers and issue Epsilon commands during the execution of a concurrent process.

With a numeric argument, the `start-process` command will create an additional concurrent process (in versions of Epsilon that support this). The `stop-process` command on `Ctrl-c Ctrl-c` will stop a running program, just as `Ctrl-c` would outside of Epsilon. You may generate an end-of-file for a program reading from the standard input by inserting a Control-Z character (quoted with `Ctrl-q`) on a line by itself, at the end of the buffer. (Use `Ctrl-q Ctrl-d` `<Enter>` for Unix.)

Programs invoked with this command should not do any cursor positioning or graphics. We provide the concurrent process facility primarily to let you run programs like compilers, linkers, assemblers, filters, etc.

`stop-process` `Ctrl-c Ctrl-c` Abort the concurrent process.

This command makes a concurrent process (see `start-process`) believe that you typed `Ctrl-c` (or, for Unix, the interrupt key).

Under Windows, providing a numeric argument makes this command send a `Ctrl-Break` signal instead of the usual `Ctrl-c` signal. A numeric argument has no effect under Unix.

`string-search` Start a search in non-incremental mode.

This command starts a search in non-incremental mode. It works like starting an incremental search with the `incremental-search` command, then disabling incremental mode with `Ctrl-o`.

`suspend-epsilon` Suspend or minimize Epsilon for Unix.

This command suspends Epsilon for Unix, returning control to the shell that launched it. Use the shell’s `fg` command to resume Epsilon. When Epsilon runs as an X11 program, it instead minimizes Epsilon’s window.

`switch-buffers` `Ctrl-⟨Tab⟩` Switch to another buffer.

This command switches to the buffer you last used. If you press `⟨Tab⟩` again while still holding down `Ctrl`, you can switch to still older buffers. Hold down `Shift` as well as `Ctrl` to move in the reverse order. You can press `Ctrl-g` to abort and return to the original buffer.

See the `switch-buffers-options` variable to customize the order in which buffers appear.

switch-windows	Switch to the next or previous window.	
This command switches to the next window. Hold down shift while pressing its key, and it will switch to the previous window.		
tabify-buffer	Replace spaces in buffer with the right number of tabs.	
This command removes all sequences of spaces and tabs throughout the buffer. In their place, it inserts a sequence of tabs followed by a sequence of spaces to reach the same column that the prior whitespace did.		
tabify-region	Ctrl-x Ctrl-Alt-i	Convert whitespace to tabs.
Between point and mark, this command removes all sequences of spaces and tabs. In their place, it inserts a sequence of tabs followed by a sequence of spaces to reach the same column that the prior whitespace did.		
tag-files	Ctrl-x Alt-.	Locate all tags in the given files.
This command prompts for a file name or file pattern. In each file, it locates each subroutine or function and makes a tag for it, so commands like goto-tag can find it later. You can use extended file patterns to tag files in multiple directories, even using a pattern like <code>**.{c,cpp,h}</code> to tag all the .c, .cpp, and .h files in a hierarchy. See page 138.		
With a prefix numeric argument, this command tags function declarations as well as function definitions, and external variable declarations as well as variable definitions. Use a numeric argument if you have an <code>#include</code> file for a package but no source file, and you want tag references to a function in the package to go to the <code>#include</code> file.		
tcl-indent-and-show-matching-delimiter	Insert braces in Tcl mode.	
The brace keys run this command in Tcl mode. It self-inserts, indents, then displays its match.		
tcl-indent-cmd	Indent this line for Tcl.	
In a line's indentation, reindent the line correctly for Tcl code. Inside the text of a line, or when repeated, insert a tab.		
If a region is highlighted, Epsilon indents all lines in the region by one tab stop. With a numeric prefix argument, Epsilon indents by that amount.		
tcl-mode	Set up for editing Tcl files.	
This command sets up syntax highlighting, indenting and other functions suitable for Tcl documents.		
telnet	Connect to a remote computer and run a shell.	
The telnet command lets you connect to a command shell on a remote computer. It puts you in a buffer that works much like the Epsilon process buffer, except the commands you type are executed on the remote machine. Provide a numeric prefix argument, or use the syntax <code>hostname:port</code> for the host name, and telnet will connect on the specified port instead of the default port. You can either use the telnet command directly, or specify a telnet: URL to find-file. (Epsilon ignores any username or password included in the URL.)		

telnet-mode Connect to a remote computer and send commands.

In Telnet mode, the key Ctrl-c Ctrl-c immediately sends an interrupt signal to the remote machine, and Ctrl-o immediately sends a Ctrl-o character (which typically makes the remote machine discard pending output). The key `<Enter>` sends a line of input to the server, while Ctrl-u `<Enter>` (running the `<Enter>` key with a numeric argument) sends a partial line.

tex-boldface TeX mode: Alt-Shift-B Make boldface text in TeX mode.

This command inserts the TeX command to make a section of text bold. You can highlight a block of text first and Epsilon will make the text bold, or you can use the command and then type the text to be bold.

tex-center-line TeX mode: Alt-s Create a centered line of text in TeX mode.

This command inserts the TeX or LaTeX command to center a line of text. (See the variable `tex-force-latex`.)

tex-close-environment TeX mode: Alt-Shift-Z Insert an `\end` for the last `\begin`.

This command searches backwards for the last `\begin{env}` directive without a matching `\end{env}` directive. Then it inserts the correct `\end{env}` directive at point.

tex-display-math TeX mode: \[Insert `\]` when you type `\[`.

When you type `\[`, this command inserts `\]` for you.

tex-environment TeX mode: Alt-Shift-E Create the specified LaTeX environment.

This command prompts for the name of a LaTeX environment, then inserts LaTeX `\begin{env}` and `\end{env}` commands for that environment. You can highlight a block of text first and Epsilon will put the environment commands around it, or you can run this command and then type the text to go in that environment. Press `?` to select an environment from a list. (The list of environments comes from the file `latex.env`, which you can edit.)

tex-fill-paragraph TeX mode: Alt-q Fill the current paragraph.

This command fills the current paragraph using the `fill-paragraph` command. But within a TeX comment, it instead fills the comment using the `fill-comment` command.

tex-footnote TeX mode: Alt-Shift-F Make a footnote in TeX mode.

This command inserts the TeX command to mark a section of text as a footnote. You can highlight a block of text first and Epsilon will make it a footnote, or you can use the command and then type the footnote.

tex-force-quote TeX mode: Alt-” Insert a " character.

This command inserts a true " character. Normally typing " itself inserts either a ‘ ‘ or a ’ ’ sequence.

tex-inline-math TeX mode: \ (Insert \) when you type \ (.

When you type \ (, this command inserts \) for you.

tex-italic TeX mode: Alt-i Make italic text in TeX mode.

This command inserts the TeX command to make a section of text italic. You can highlight a block of text first and Epsilon will make the text italic, or you can use the command and then type the italic text.

tex-left-brace TeX mode: { Insert } when you type {.

This command inserts a matched pair of braces. After a \ character, it inserts a \ before the closing brace. But if you type this key just before a non-whitespace character, it inserts only a {. This makes it easier to surround existing text with braces.

tex-math-escape TeX mode: \$ Insert \$ when you type \$.

This command inserts a matched pair of \$ characters (except after a \ character).

tex-mode Set up for editing TeX or LaTeX documents.

This command sets up Epsilon for editing TeX or LaTeX documents. Keys in TeX mode include Alt-i for italic text, Alt-Shift-I for slanted text, Alt-Shift-T for typewriter, Alt-Shift-B for boldface, Alt-Shift-C for small caps, Alt-Shift-F for a footnote, and Alt-s for a centered line. Alt-Shift-E prompts for the name of a LaTeX environment, then inserts \begin{env} and \end{env} lines.

For all these commands, you can highlight a block of text first and Epsilon will make the text italic, slanted, etc. or you can use the command and then type the text to be italic, slanted, etc.

The keys ‘{’ and ‘\$’ insert matched pairs of characters (either {} or \$\$), the keys ⟨Comma⟩ and ⟨Period⟩ remove a preceding italic correction \/, the " key inserts the appropriate kind of doublequote sequence like ‘ ‘ or ’ ’, and Alt-" inserts an actual " character.

tex-mouse-select TeX mode: M-⟨Left⟩ With Ctrl, show TeX output.

In a TeX buffer, you can use the jump-to-dvi command to make your DVI previewer or PDF viewer display the output resulting from the text near point. See page 98 for instructions on setting this up. This command lets you run jump-to-dvi by left-clicking while holding down the Ctrl key. Without Ctrl it acts like a normal left-click.

tex-quote TeX mode: ” Insert the right TeX doublequote sequence.

This command inserts the appropriate doublequote sequence like ‘ ‘ or ’ ’, based on the preceding characters. Alt-" inserts an actual " character.

tex-rm-correction TeX mode: ⟨Comma⟩, ⟨Dot⟩ Remove an italic correction.

This command removes any nearby italic correction \/ when appropriate.

`tex-slant` TeX mode: Alt-Shift-I Make slanted text in TeX mode.

This command inserts the TeX command to make a section of text slanted. You can highlight a block of text first and Epsilon will make the text slanted, or you can use the command and then type the text to be slanted.

`tex-small-caps` TeX mode: Alt-Shift-C Make small caps text in TeX mode.

This command inserts the TeX command to set a section of text in small caps. You can highlight a block of text first and Epsilon will put the text in small caps, or you can use the command and then type the text.

`tex-typewriter` TeX mode: Alt-Shift-T Use a typewriter font in TeX mode.

This command inserts the TeX command to set a section of text in a typewriter font. You can highlight a block of text first and Epsilon will set that text in a typewriter font, or you can use the command and then type the text.

`this-buffer-grep` Search in this buffer for a pattern.

This command searches through the current buffer for a pattern. It prompts for a search string, then lists all matching lines in the grep buffer. The grep buffer then appears in the current window.

By default, `this-buffer-grep` interprets the search string as a regular expression. Press Ctrl-t at the search string prompt to toggle regular expression mode. You can also type Ctrl-w or Ctrl-c to toggle word-mode or case-folding searches, respectively.

In grep mode, alphabetic keys run special grep commands. See the description of the `grep-mode` command for details. Typing H or '?' in grep mode gives help on grep subcommands.

A bit in the `search-in-region` variable makes grepping a single buffer restrict itself to a highlighted region. By default, Epsilon searches the entire buffer.

`to-indentation` Alt-m Move point to the end of the indentation.

This command positions point before the first non-whitespace character in the line.

`to-left-edge` Brief: Shift-(Home) Move to the left edge of the window.

This command moves point to the left edge of the current window.

`to-right-edge` Brief: Shift-(End) Move to the right edge of the window.

This command moves point to the right edge of the current window.

`toggle-borders` Brief: Alt-F1 Remove borders around windows,
use color to distinguish them.

This command removes the borders around ordinary tiled windows, letting the text regions occupy more of the screen. If the windows have no borders already, this command restores them. When this command reenables borders, it does so according to the settings of the variables `border-left`, `border-top`, and so forth. Epsilon displays a border only if the appropriate variable has been set, and `toggle-borders` hasn't disabled all borders.

When there are no window borders, Epsilon provides each window with its own separate color scheme, in place of the single one selected by `set-color`. (You can still use `set-color` to set the individual colors in a color scheme, but Epsilon doesn't care which particular color scheme you select when it displays the contents of individual windows. It does use the selected color scheme for other parts of the screen like the echo area or screen border.)

The color schemes Epsilon uses for borderless windows have names like "window-black", "window-blue" and so forth. Epsilon assigns them to windows in the same order they appear in `set-color`. You can remove one from consideration using the `delete-name` command, or create a new one using `set-color` (give it a name starting with "window-").

`toggle-case-fold` Turn off case folding in search or turn it on again.

This command toggles the state of case folding for the current buffer by setting the `case-fold` variable. When case folding is on, searching and other operations ignore the case of letters. With a nonzero numeric argument, the command enables case folding instead of toggling it; with a zero numeric argument, it turns off case folding.

`toggle-menu-bar` Toggle whether a permanent menu bar appears.

Add a menu bar at the top of the screen, moving windows down one line. If Epsilon already displays a menu bar, remove it.

`toggle-scroll-bar` Toggle whether tiled windows have permanent scroll bars.

Put a scroll bar on the right edge of all tiled windows. If tiled windows already have scroll bars, remove them.

`toggle-toolbar` Turn the tool bar on or off.

The Windows GUI versions of Epsilon can display a tool bar. Position the mouse over a tool bar button for a moment and Epsilon will describe what it does. This command hides or displays the tool bar.

`transpose-characters` Ctrl-t Swap the characters around point.

At the end of a line, the command switches the two previous characters. At the beginning of a line, it switches the following two characters. Otherwise, it switches the characters before and after point. If the current line has less than two characters, however, nothing happens. Point never changes.

`transpose-lines` Ctrl-x Ctrl-t Swap the current and previous lines.

After the exchange, the command positions point between the two lines.

`transpose-words` Alt-t Swap the current and previous words.

The command leaves untouched the text between the words. After the exchange, the command positions point between the two words.

`tutorial` This command shows Epsilon's tutorial.

unbind-key Remove the binding from a key.

The unbind-key command prompts for a key and then offers to rebind the key to the normal-character command, or to remove any binding it may have. A key bound to normal-character will self-insert; that's how keys like 'j' are bound. A key with no binding at all simply displays an error message.

undo F9 Undo the last buffer change or movement.

This command undoes the last change you made to the buffer. If repeated, it undoes earlier changes. You may reinstate the changes with redo.

undo-by-commands Undo, grouping by commands.

This command operates like undo, except that it groups editing operations based on commands, similar to Brief.

undo-changes Ctrl-F9 Undo, grouping movements.

This command operates like undo, except that it will automatically undo any run of changes to the buffer that involve only movements of point as a unit, and stop just before a change of actual buffer contents. When you invoke undo-changes, it performs an undo, and if the undone operation simply moved point, then it continues to undo changes until it encounters a non-movement operation.

undo-movements Ctrl-F11 Return to the site of previous edits.

After making changes at various spots in a buffer, use this command and redo-movements to move to the location of previous edits. Epsilon uses undo information to record where in a buffer editing has occurred. This command moves backward through the list of locations, toward the oldest location.

unicode-convert-from-encoding Interpret buffer as Unicode encoding.

This command converts multi-character sequences in the buffer, representing individual Unicode characters, into the Unicode characters they represent. It prompts for the name of the encoding scheme it should use.

If there's a highlighted region, only characters in the region are converted in this way.

unicode-convert-to-ascii Substitute characters similar to ASCII ones.

This command replaces some Unicode and Windows-1252 punctuation characters in the buffer with their nearest ASCII equivalent. It converts various quote-like characters to ' or ", various hyphen-like characters to -, and so forth, though it does not try to strip accents from letters.

If there's a highlighted region, only characters in the region are converted in this way.

unicode-convert-to-encoding Replace Unicode characters with encoded versions.

This command prompts for the name of an encoding, then converts each character in the buffer to its representation in that encoding.

If there's a highlighted region, only characters in the region are converted in this way.

`uniq` Remove extra copies of duplicate lines.

The command goes through the current buffer and looks for adjacent identical lines, deleting the duplicate copies of each repeated line and leaving just one. It doesn't modify any lines that only occur once. If the `case-fold` variable is nonzero, lines that only differ by case will be considered identical. Also see the `keep-unique-lines` and `keep-duplicate-lines` command.

`untabify-buffer` Replace tabs in the buffer with spaces.

This command replaces each tab in the buffer by the number of spaces required to fill the same number of columns.

`untabify-region` `Ctrl-x Alt-i` Convert tabs to spaces between point and mark.

This command replaces each tab between point and mark by the number of spaces required to fill the same number of columns.

`untag-files` Discard tags for one or more files.

This command constructs a list of all files represented in the current tag file. You can edit the list in a recursive edit. When you exit the recursive edit with the `exit-level` command on `Ctrl-x Ctrl-z`, any files you've removed from the list will be untagged.

`up-line` `Ctrl-p` Point moves to the previous line.

The command tries to keep point near the same horizontal position.

`uppercase-word` `Alt-u` Make the current word upper case.

Point travels forward through the buffer as with `forward-word`, changing all the letters it encounters to upper case. If the current buffer contains a highlighted region, `Epsilon` instead changes all the letters in the region to upper case, leaving point unchanged.

`vbasic-mode` Set up for editing Visual Basic.

This command puts the current buffer in a mode suitable for editing Visual Basic or similar languages (like VBscript or VBA). Syntax highlighting, indenting, tagging, delimiter highlighting and commenting commands are all provided.

`vhdl-indenter` `<Tab>` in VHDL mode Indent this line for VHDL.

In a line's indentation, reindent the line correctly for VHDL code. Inside the text of a line, or when repeated, insert a tab.

If a region is highlighted, `Epsilon` indents all lines in the region by one tab stop. With a numeric prefix argument, `Epsilon` indents by that amount.

`vhdl-mode` Set up for editing VHDL files.

This command puts the current buffer in a mode suitable for editing VHDL files.

view-lugaru-web-site Connect to Lugaru's web site.

This command starts your web browser and points it to Lugaru's web site. It only works under Epsilon for Windows on systems with more recent web browsers, and in Epsilon for Unix under X11.

view-process Shift-F3 Pop up a window of process output;
pick an error msg.

This command pops up a window showing the process buffer, including all compiler command lines and any resulting error messages. You can move to any line and press `<Enter>`, and Epsilon will immediately locate the error message on the current line (or a following line) and move to the file and line number in error.

view-web-site Pass a URL to a browser.

This command prompts for a URL, scanning the current buffer for a suitable default. Then it starts your web browser and passes the URL to it. In Epsilon for Unix, it only works under X11.

visit-file Ctrl-x Ctrl-v Read a file into the current buffer.

This command prompts for a file name, then reads that file into the current buffer, and positions point to the beginning. If no file with the given name exists, it creates a blank buffer. In either case, the command discards the old buffer contents.

Before discarding modified buffers, the command asks if you want to save the current buffer contents. With a numeric argument, it asks no questions. This comes in handy for reverting the buffer to the contents of its file.

visual-diff Use color-coding to compare two buffers.

The visual-diff command is like the diff command but uses colors to show differences. It compares the current buffer with the one shown in the next window on the screen, and constructs a new buffer that contains all the lines of the two buffers. Lines from the first buffer that don't appear in the second are displayed with a red background. Lines in the second buffer that don't appear in the first have a yellow background. Lines that are the same in both buffers are colored normally, as are those parts of an added or removed line that are the same in both buffers. With a numeric argument, the command prompts for the name of a buffer to hold the results, instead of using a buffer named `#diff#`.

visual-diff-mode Use color-coding to compare two buffers.

The visual-diff command creates a buffer in visual diff mode that shows the changes between one buffer and another. Added lines are shown with a yellow background, deleted lines are shown with a red background, and common lines are colored as in the original buffers.

In a visual-diff buffer, the keys `Alt-⟨Down⟩` and `Alt-]` move to the start of the next changed or common section. The keys `Alt-⟨Up⟩` and `Alt-[` move to the previous one.

wall-chart Make a chart of the current key bindings.

This command creates a wall chart consisting of all bound keys and their current bindings. You can print it using the `print-buffer` command.

web-search Search the web for some text.

This command performs a web search for the word at the cursor, or the highlighted text if any, displaying the result in your web browser. It appends the given text to the `web-search-url` variable, which you can modify to choose a different search provider. The right-button context menu in Epsilon for Windows runs this command when you choose Web Search.

what-is F6, F1 Q Find a command bound to a key.

The command prompts for a key, then displays a message telling what command runs when you press that key.

widen-buffer Restore normal access to the current buffer.

This command gives you normal access to the buffer. Use it after a `narrow-to-region` command to cancel the effect of that command.

write-file Ctrl-x Ctrl-w Write the buffer to a file.

This command prompts for a file name, then writes the buffer to a file with that name. The file associated with the current buffer becomes that file, so subsequent uses of the `save-file` command will write the buffer to that file without asking for a file name. See also `copy-to-file` and `save-file`. As usual, typing Alt-g at the prompt inserts the current file name for editing.

Epsilon writes the file using the same line termination rules and encoding as it used when it read the file, perhaps altered by a subsequent use of the `set-line-translate` or `set-encoding` commands. With a numeric prefix argument, the `write-file` command first asks for a new translation rule and encoding.

write-files-and-exit Brief: Ctrl-x Save modified files, then leave Epsilon.

This command saves all modified buffers except those that do not have files associated with them. If there are no errors, it then exits Epsilon.

write-region Ctrl-x w Write the region to the specified file.

The command prompts for a file name, then writes the characters between point and mark to that file.

write-session Record the current file & window configuration.

The new `write-session` command writes a session file, detailing the files you're currently editing, the window configuration, default search strings, and so forth. By default, Epsilon writes a session file automatically whenever you exit, but you can use this command if you prefer to save and restore sessions manually.

write-state Ctrl-F3 Save all commands and variables for later automatic loading.

This command prompts for a file name. It alters any extension to ".sta", and then loads the documentation file and records the position of each of the definitions in it (to speed up the help system). Epsilon then writes all its commands, variables, and bindings to the named file. Restarting Epsilon with the command "`epsilon -sfilename`", where "filename" denotes the name of the state file, makes Epsilon use the commands in that file. Epsilon normally uses the state file "`epsilon-v14.sta`".

xml-mode Set up for editing XML files.

This command puts the current buffer in XML mode. Epsilon will do syntax-highlighting, smart indenting, and brace-matching.

xml-sort-by-attribute-name XML mode: Alt-Shift-R Reorder attributes, align them.

In XML and HTML modes, this command sorts the attributes of the current tag alphabetically by attribute name. With a highlighted region, it sorts the attributes of each tag in the region, then aligns corresponding attributes so they start at the same column in each tag. It uses the first tag in the region for its list of attribute names to be aligned, so attributes that aren't used in the first tag will not be aligned.

yank Ctrl-y Insert the contents of a kill buffer.

This command inserts the contents of the last kill buffer at point, then positions point after the insertion, and the mark before it. In some modes this command then reindents the inserted text. See the `reindent-after-yank` variable. If another program has placed text on the system clipboard, this command will use it instead of the kill buffer, except in keyboard macros. See the `clipboard-access` variable for more information.

If the kill buffer contains a rectangle, the command inserts it at the current column, on the current and successive lines. It shifts existing text to the right, unless you've enabled overwrite mode, in which case the block replaces any existing text in those columns. (When yanked text comes from the system clipboard, Epsilon never treats it as a rectangle.)

yank-pop Alt-y Cycle through previous kill buffers.

This command replaces the just-yanked kill buffer with the contents of the previous kill buffer. It only works after a `yank` or `yank-pop` command.

yank-x-selection Insert selection from X11 Window System.


When Epsilon for Unix runs under the X11 window system, yanking normally inserts text from its clipboard, and killing text normally puts it on the clipboard as well. But X11 has more than one way of transferring text from one program into another. An older method uses the "primary selection". Selecting text with the mouse sets the primary selection, and middle-clicking the mouse (with Shift, or after setting the `mouse-center-yanks` variable to 1) pastes it. This command offers another way to paste X11's primary selection.

zoom-window Ctrl-x z Zoom in on the current window.

This command, like the `one-window` command, makes the current window occupy the entire screen. But it also saves away the old window configuration. Later, when you invoke `zoom-window` again, it restores the old window configuration.

Chapter 6

Variables



This chapter lists all of Epsilon's user-settable variables, with the exception of some variables used only within a particular subsystem, and not meant to be set by the user.

The variables typically modified to customize Epsilon are marked "Preference". Be careful when setting variables marked "System". They should generally be set only via the appropriate EEL commands, not directly by the user. By default, the `set-variable` and `edit-variables` commands omit system variables.

`abort-file-io` System Default: -3

If the user presses the abort key when Epsilon is reading or writing a file, Epsilon asks whether to abort the file input/output operation. If the user says to abort, the file reading or writing primitive returns an error code, `EREADABORT` or `EWRITEABORT`, respectively. This variable changes Epsilon's behavior; EEL code may set it using a `save_var` statement to one of these values: A value of `ABORT_IGNORE` (0) makes Epsilon ignore the abort key during file I/O. A value of `ABORT_ASK` (-3, the default) makes Epsilon ask the user as above and return an error code. A value of `ABORT_ERROR` (-2) make Epsilon return the error code without prompting first. The value `ABORT_JUMP` (-1) doesn't ask either; it calls the `check_abort()` primitive, which typically aborts the entire command.

`abort-file-matching` Default: 0

Epsilon's file matching primitives respond to the abort key based on the value of this variable. If 0, they ignore the abort key. If 1, they abort out of the calling function. If 2, they return an error code. EEL functions that are prepared to handle aborting should set this variable.

`abort-key` System Default: 7 (Ctrl-g)

Epsilon aborts the current command when you press the key whose value is `abort-key`. To disable the abort key, set `abort-key` to -1. By default, the `abort-key` variable is set to Control-G. For correct behavior, use the `set-abort-key` command to set this variable.

`abort-searching` Default: -1

If the user presses the abort key during searching, Epsilon's behavior depends upon the value of the `abort-searching` variable. If it's 0, the key is ignored and the search continues. If it's `ABORT_JUMP` (-1, the default), Epsilon aborts the search and jumps by calling the `check_abort()` primitive. If it's `ABORT_ERROR` (-2), Epsilon aborts the search and returns the value `ABORT_ERROR`. The `search()`, `re_search()`, `re_match()`, and `buffer_sort()` primitives all use the `abort-searching` variable to control aborting.

`align-region-extra-space` Preference Default: 0

You can use this variable to increase the space between columns aligned using the `align-region` command. To have `align-region` apply additional space just once, run it with a numeric prefix argument.

`align-region-rules` Preference Default: 1023

The `align-region` command uses a series of mode-specific alignment rules. Each of these rules corresponds to a bit in this variable that enables using that rule. Some rules only apply in certain modes. The rules: align the first = on a line (1, default rule, most modes), align comments (2, default rule, most modes), align variable names being defined (4, C mode), align definitions with `#define` (8, C mode), align a backslash character at the end of a line (16, C mode).

`all-must-build-mode` Default: 0

Epsilon “precomputes” most of the text of each mode line, so it doesn’t have to figure out what to write each time it updates the screen. Setting the `all-must-build-mode` variable nonzero warns Epsilon that all mode lines must be rebuilt. Epsilon resets the variable to zero after every screen update.

`all-read-only` Default: 0

While this variable is non-zero, every time Epsilon reads a file it will set its buffer to read-only mode.

`already-made-backup` System Buffer-specific Default: 0

Epsilon sets this buffer-specific variable nonzero whenever it saves a file and makes a backup.

`alt-invokes-menu` Preference Default: 0

In a typical Windows program, pressing and releasing the Alt key without pressing any other key moves to the menu bar, highlighting its first entry. Set this variable to 1 if you want Epsilon to do this. The variable has no effect on what happens when you press Alt and then press another key before releasing Alt: this will run whatever command is bound to that key. If you want Alt-e, for example, to display the Edit menu, you can bind the command `invoke-windows-menu` to it.

`alt-numpad-keys` Preference Default: 0

Epsilon for Windows normally interprets Alt plus Numeric Pad keys as a way to enter a character by its numeric code, as in other Windows programs. Type a number on the numeric pad while holding down Alt to use this feature.

But if this variable is nonzero, Epsilon instead treats Alt plus Numpad keys as individual keys, which can be bound to commands like any other keys.

`anon-ftp-password` Preference Default:
"-EpsilonUser@unknown.host"

When Epsilon uses FTP to read or write files to a computer on the Internet, and logs in anonymously, it provides the contents of this variable as a password. (Anonymous FTP sites ask that you provide your email address as a password when you log in anonymously.) You can set this to your email address.

`argc` System Default: varies

The `argc` variable contains the number of words on Epsilon’s command line, after Epsilon removes several flags it processes internally. The count includes the command name “epsilon” at the start of the command line.

`auto-fill-comment-rules` Preference Buffer-specific Default: 3

Bits in this variable let you customize how Epsilon breaks and fills comment lines, in modes where Epsilon does this. The value 1 lets Epsilon break an overlong comment line when you press `<Enter>`. The value 2 lets Epsilon break an overlong comment line when you insert any other character. The value 4 applies to C-style block comments only, when Epsilon continues a comment line following one that starts with `/*`. It tells Epsilon to create new lines with a prefix of a single space, not a space, star, and space. The value 8 tells Epsilon not to merge adjacent block comments when filling, as it does by default.

`auto-fill-indents` Preference Buffer-specific Default: 1

When Epsilon automatically inserts new lines for you in auto fill mode, it indents new lines (by calling the indenter function for the current buffer) only if the buffer-specific variable `auto-fill-indents` has a nonzero value.

`auto-indent` Preference Buffer-specific Default: 0

Epsilon can automatically indent for you when you press `<Enter>`. Setting the buffer-specific variable `auto-indent` nonzero makes Epsilon do this. The way Epsilon indents depends on the current mode. For example, C mode knows how to indent for C programs. In Epsilon's default mode, fundamental mode, Epsilon indents like `indent-previous` if you set `auto-indent` nonzero.

`auto-menu-bar` Preference Default: 1

If nonzero, moving the mouse past the top edge of the screen makes Epsilon for DOS display the menu bar, in environments that support this. Other versions of Epsilon ignore this variable.

`auto-read-changed-file` Preference Buffer-specific Default: 0

If nonzero, when Epsilon notices that a file on disk has a different timestamp than the file in memory, it automatically reads the new version of the file and displays a message to that effect. Epsilon won't do this if you've edited the copy of the file in memory, or if the file's disk size is substantially smaller than it was. In those cases, Epsilon asks what to do. Also see the variable `want-warn`.

`auto-save-biggest-file` Preference Default: 5,000,000

Buffers larger than this many characters will never be auto-saved.

`auto-save-count` Preference Default: 500

When `want-auto-save` is nonzero, Epsilon automatically saves a copy of each unsaved file every `auto-save-count` keystrokes.

`auto-save-idle-seconds` Preference Default: 30

When `want-auto-save` is nonzero, Epsilon automatically saves a copy of each unsaved file after you've been idle for this many seconds.

`auto-save-name` Preference Default: "%p#%b%e.asv#"

When `want-auto-save` is nonzero, Epsilon regularly saves a copy of each unsaved file. This variable contains a template which determines how Epsilon chooses the file name for the autosaved file. Epsilon substitutes pieces of the original file name for codes in the template, as follows (examples are for the file `c:\dos\read.me`):

- %p** The original file's path (`c:\dos\`).
- %b** The base part of the original file name (`read`).
- %e** The extension of the original file name (`.me`).
- %f** The full name of the original file (`c:\dos\read.me`).
- %r** The name of the file relative to the current directory. (`read.me` if the current directory is `c:\dos`, `dos\read.me` if the current directory is `c:\`, otherwise `c:\dos\read.me`).
- %x** The full pathname of the directory containing the Epsilon executable.
- %X** The full pathname of the directory containing the Epsilon executable, after converting all Windows long file names to their equivalent short name aliases.

By default, Epsilon writes to a file in the same directory but with a `#` character before the name and `.asv#` after it. If you change this to `%f` (to auto-save to the original file name) or some other setting that could name a file you want to keep, be sure to set appropriate bits in the `want-auto-save` variable to disable automatic deletion of auto-save files.

`auto-save-tags` Preference Default: 1

Whenever Epsilon retags a file, it saves the modified tag file. Set this variable to zero to disable this behavior; you can then save the tag file as you would any other file.

`auto-show-adjacent-delimiter` Preference Default: 3

When the cursor is on a delimiter character in various language modes, Epsilon highlights the character and its match. Epsilon can also highlight when the cursor is adjacent to a delimiter character. Bits in this variable control this. The 1 bit makes Epsilon highlight if the cursor is just past a right-hand delimiter. The 2 bit makes Epsilon highlight if the cursor is just past a left-hand delimiter. The 4 bit lets highlighting appear in all windows showing the current buffer, not just the current one. (Changing this last setting won't affect existing buffers.)

`auto-show-batch-delimiters` Preference Default: 1

When the cursor is next to a brace, bracket, or parenthesis in Batch file mode, Epsilon will try to locate its matching brace, bracket, or parenthesis, and highlight them both. If the current character has no match, Epsilon will not highlight it. Set this variable to zero to disable this feature. Also see the `auto-show-adjacent-delimiter` variable.

`auto-show-c-delimiters` Preference Default: 3

When the cursor is next to a brace, bracket, or parenthesis in C mode, Epsilon will try to locate its matching brace, bracket, or parenthesis, and highlight them both. If the current character has no match, Epsilon will not highlight it. This variable's 1 bit enable that. Similarly, Epsilon will highlight matching `#if/#else/#endif` preprocessor lines. This variable's 2 bit enables that, while its 4 bit makes Epsilon skip highlighting when point is after the preprocessor keyword, such as in the expression part of an `#if`. Also see the `auto-show-adjacent-delimiter` variable.

`auto-show-conf-delimiters` Preference Default: 1

When the cursor is next to a brace, bracket, or parenthesis in Conf file mode, Epsilon will try to locate its matching brace, bracket, or parenthesis, and highlight them both. If the current character has no match, Epsilon will not highlight it. Set this variable to zero to disable this feature. Also see the `auto-show-adjacent-delimiter` variable.

`auto-show-css-delimiters` Preference Default: 1

When the cursor is next to a brace, bracket, or parenthesis in CSS mode, or in a CSS block within an HTML file, Epsilon will try to locate its matching brace, bracket, or parenthesis, and highlight them both. If the current character has no match, Epsilon will not highlight it. Set this variable to zero to disable this feature. Also see the `auto-show-adjacent-delimiter` variable.

`auto-show-delimiter-delay` System Default: 5

Epsilon uses this variable internally to decide how long to wait before searching and highlighting matching delimiters.

`auto-show-gams-delimiters` Preference Default: 1

When the cursor is next to a bracket or parenthesis in GAMS mode, Epsilon will try to locate its matching bracket or parenthesis, and highlight them both. If the current character has no match, Epsilon will not highlight it. Set this variable to zero to disable this feature. Also see the `auto-show-adjacent-delimiter` variable.

`auto-show-html-delimiters` Preference Default: 7

Bits in this variable control automatic highlighting in HTML, XML, and PHP modes. When the cursor is next to a `<` or `>` character in HTML mode, Epsilon will try to locate its matching `>` or `<` and highlight them both. The 1 bit enables this. If the current character has no match, Epsilon will not highlight it. Also see the `auto-show-adjacent-delimiter` variable.

The 2 bit tells Epsilon to look for matching start and end tags. Inside such tag names, Epsilon will highlight both the start and end tags. Tags with no match will not be highlighted. The 4 bit controls how Epsilon highlights tags that have no match but should. If on, they receive a different color; otherwise they're not highlighted.

`auto-show-matching-characters` System Buffer-specific Default: none

Epsilon's `auto-show-delimiters` feature stores the set of delimiter characters for the current mode in this variable.

`auto-show-nsis-delimiters` Preference Default: 3

When the cursor is next to a brace, bracket, or parenthesis in NSIS mode, Epsilon will try to locate its matching brace, bracket, or parenthesis, and highlight them both. If the current character has no match, Epsilon will not highlight it. Set this variable to zero to disable this feature. Also see the `auto-show-adjacent-delimiter` variable.

`auto-show-perl-delimiters` Preference Default: 1

When the cursor is next to a brace, bracket, or parenthesis in Perl mode, Epsilon will try to locate its matching brace, bracket, or parenthesis, and highlight them both. If the current character has no match, Epsilon will not highlight it. Set this variable to zero to disable this feature. Also see the `auto-show-adjacent-delimiter` variable.

`auto-show-php-delimiters` Preference Default: 1

When the cursor is next to a brace, bracket, or parenthesis in PHP mode, Epsilon will try to locate its matching brace, bracket, or parenthesis, and highlight them both. If the current character has no match, Epsilon will not highlight it. Set this variable to zero to disable this feature. Also see the `auto-show-adjacent-delimiter` variable.

`auto-show-postscript-delimiters` Preference Default: 1

When the cursor is next to a bracket or parenthesis in PostScript mode, Epsilon will try to locate its matching brace, bracket or parenthesis, and highlight them both. If the current character has no match, Epsilon will not highlight it. Set this variable to zero to disable this feature. Also see the `auto-show-adjacent-delimiter` variable.

`auto-show-python-delimiters` Preference Default: 1

When the cursor is next to a brace, bracket, or parenthesis in Python mode, Epsilon will try to locate its matching brace, bracket or parenthesis, and highlight them both. If the current character has no match, Epsilon will not highlight it. Set this variable to zero to disable this feature. Also see the `auto-show-adjacent-delimiter` variable.

`auto-show-rmode-delimiters` Preference Default: 1

When the cursor is next to a brace, bracket, or parenthesis in R mode, Epsilon will try to locate its matching brace, bracket, or parenthesis, and highlight them both. If the current character has no match, Epsilon will not highlight it. Set this variable to zero to disable this feature. Also see the `auto-show-adjacent-delimiter` variable.

`auto-show-shell-delimiters` Preference Default: 1

When the cursor is next to a brace, bracket, or parenthesis in Shell mode, Epsilon will try to locate its matching brace, bracket, or parenthesis, and highlight them both. If the current character has no match, Epsilon will not highlight it. Set this variable to zero to disable this feature. Also see the `auto-show-adjacent-delimiter` variable.

`auto-show-tcl-delimiters` Preference Default: 1

When the cursor is next to a brace, bracket, or parenthesis in Tcl mode, Epsilon will try to locate its matching brace, bracket, or parenthesis, and highlight them both. If the current character has no match, Epsilon will not highlight it. Set this variable to zero to disable this feature. Also see the `auto-show-adjacent-delimiter` variable.

`auto-show-tex-delimiters` Preference Default: 1

When the cursor is next to a curly brace or square bracket character like {, }, [, or] in TeX mode, Epsilon will try to locate its matching character and highlight them both. If the current character has no match, Epsilon will not highlight it. Set this variable to zero to disable this feature. Also see the `auto-show-adjacent-delimiter` variable.

`auto-show-vbasic-delimiters` Preference Default: 1

When the cursor is next to a brace, bracket, or parenthesis in Visual Basic mode, Epsilon will try to locate its matching brace, bracket or parenthesis, and highlight them both. If the current character has no match, Epsilon will not highlight it. Set this variable to zero to disable this feature. Also see the `auto-show-adjacent-delimiter` variable.

`auto-show-vhdl-delimiters` Preference Default: 1

When the cursor is next to a parenthesis in VHDL mode, Epsilon will try to locate its matching parenthesis, and highlight them both. If the current character has no match, Epsilon will not highlight it. Set this variable to zero to disable this feature. Also see the `auto-show-adjacent-delimiter` variable.

`avoid-bottom-lines` Preference Default: 1

This variable tells Epsilon how many screen lines at the bottom of the screen are reserved, and may not contain tiled windows. By default, this variable is one, to make room for the echo area.

`avoid-top-lines` Preference Default: 0

This variable tells Epsilon how many screen lines at the top of the screen are reserved, and may not contain tiled windows. By default, this variable is zero, indicating that tiled windows reach to the top of the screen. If you create a permanent menu bar, Epsilon sets this variable to one.

`backup-by-renaming` Preference Default: 2

When the `want-backups` variable tells Epsilon to back up files before it saves them, it can do this in one of two ways, either by renaming the original file and then writing a new file by that name, or by copying the data from the original file to the backup file, then modifying the original file.

Renaming is faster, but if the file system provides the original file with some special properties or attributes, and Epsilon doesn't know how to copy them from one file to another, the properties may become associated with the backup file, not the original. Copying instead of renaming avoids this problem.

If this variable is zero, Epsilon always copies files to back them up. If this variable is one, Epsilon tries to rename the file, only copying the file if it can't successfully rename the file. If this variable is two, the default, Epsilon copies under Unix (since some security-enhanced versions of Linux use just such attributes) and tries to rename first under Windows.

backup-name	Preference	Default: "%p%b.bak"
-------------	------------	---------------------

If you've set `want-backups nonzero`, telling Epsilon to make a backup whenever it saves a file, Epsilon uses this variable to construct the name of the backup file. The variable contains a template, which Epsilon copies, substituting pieces of the original file for codes in the template, as follows (examples are for the file `c:\dos\read.me`):

- %p** The original file's path (`c:\dos\`).
- %b** The base part of the original file name (`read`).
- %e** The extension of the original file name (`.me`).
- %f** The full name of the original file (`c:\dos\read.me`).
- %r** The name of the file relative to the current directory. (`read.me` if the current directory is `c:\dos`, `dos\read.me` if the current directory is `c:\`, otherwise `c:\dos\read.me`).
- %x** The full pathname of the directory containing the Epsilon executable.
- %X** The full pathname of the directory containing the Epsilon executable, after converting all Windows long file names to their equivalent short name aliases.

By default, Epsilon renames the old file so it has extension `".bak"`.

batch-auto-show-delim-chars		Default: "[()]"
-----------------------------	--	-----------------

This variable holds the set of delimiter characters that should trigger Epsilon's `auto-show-delimiters` feature in Batch mode. Epsilon will search for and highlight the match of each delimiter.

beep-duration	Preference	Default: 5
---------------	------------	------------

This variable specifies the duration of Epsilon's warning beep, in hundredths of a second. If zero, Epsilon uses a default beeping sound. Under Windows and Unix, setting the variable has no effect.

beep-frequency	Preference	Default: 370
----------------	------------	--------------

This variable specifies the frequency of Epsilon's warning beep in hertz. If zero, Epsilon instead flashes the mode line of each window for a moment. Under Windows, setting the variable has no effect. Under Unix, Epsilon will flash if the variable is zero, but won't change the frequency.

bell-on-abort	Preference	Default: 0
---------------	------------	------------

If nonzero, Epsilon will beep when you abort a command, press an unbound key, or when a command aborts due to an error.

bell-on-autosave-error	Preference	Default: 1
------------------------	------------	------------

If nonzero, Epsilon will beep when it can't autosave a file.

bell-on-bad-key	Preference	Default: 1
-----------------	------------	------------

If nonzero, Epsilon will beep when you press an illegal option at a prompt.

<code>bell-on-completion</code>	Preference	Default: 1
If nonzero, Epsilon will beep when it's completing on command names, file names, or similar things, and it can't find any matches.		
<code>bell-on-date-warning</code>	Preference	Default: 1
If nonzero, Epsilon will beep when it puts up its warning that a file has been changed on disk.		
<code>bell-on-read-error</code>	Preference	Default: 1
If nonzero, Epsilon will beep when it gets an error reading a file.		
<code>bell-on-search</code>	Preference	Default: 1
If nonzero, Epsilon will beep when it can't find the text you're searching for.		
<code>bell-on-write-error</code>	Preference	Default: 1
If nonzero, Epsilon will beep when it gets an error writing a file.		
<code>border-bottom</code>	Preference	Default: 0
If nonzero, Epsilon puts a border on the bottom edges of tiled windows that touch the bottom of the screen (or the echo area, if it's at the bottom of the screen). If Epsilon is set to display a mode line below each tiled window, it puts a border there too, regardless of this variable's setting. If you've run the <code>toggle-borders</code> command to suppress borders entirely, you must run that command again to reenable the borders.		
<code>border-inside</code>	Preference	Default: 1
If nonzero, Epsilon puts a vertical border between two side-by-side tiled windows. If you've run the <code>toggle-borders</code> command to suppress borders entirely, you must run that command again to reenable the borders.		
<code>border-left</code>	Preference	Default: 0
If nonzero, Epsilon puts a border on the left edges of tiled windows that touch the left edge of the screen. If you've run the <code>toggle-borders</code> command to suppress borders entirely, you must run that command again to reenable the borders.		
<code>border-right</code>	Preference	Default: 0
If nonzero, Epsilon puts a border on the right edges of tiled windows that touch the right edge of the screen. If you've run the <code>toggle-borders</code> command to suppress borders entirely, you must run that command again to reenable the borders.		
<code>border-top</code>	Preference	Default: 0
If nonzero, Epsilon puts a border on the top edges of tiled windows that touch the top edge of the screen. If nonzero, Epsilon puts a border on the top edges of tiled windows that touch the top of the screen (or the echo area, if it's at the top of the screen). If Epsilon is set to display a mode line above each tiled window, it puts a border there too, regardless of this variable's setting. If you've run the <code>toggle-borders</code> command to suppress borders entirely, you must run that command again to reenable the borders.		

<code>browser-file</code>	Preference	Default: ""
---------------------------	------------	-------------

The `browse-symbol` command retrieves its data from the browser database file named by this variable. Use the `select-browse-file` command to set it.

<code>browser-filter</code>	Preference	Default: 127
-----------------------------	------------	--------------

The `browse-set-filter` command saves its filter setting in this variable.

<code>browser-filter-usedby</code>	Preference	Default: 127
------------------------------------	------------	--------------

The `browse-set-usedby-filter` command saves its filter setting in this variable.

<code>browser-options</code>	Preference	Default: 0x7fff
------------------------------	------------	-----------------

Bits in this variable controls various options for the `#symbols#` buffer. It's constructed by the `browse-symbol` command, and when following tags using a browser database file.

These bits apply only when a tagging command looks up a symbol, and Epsilon generates a `#symbols#` buffer for it as a side effect:

The 0x1 bit lets Epsilon include actual source file lines in the `#symbols#` buffer. If off, the buffer will just show their file names and line numbers.

The 0x2 bit includes references to the symbol, not just definitions, when it lists source file lines. That is, it includes the section labeled "is used at".

The 0x4 bit includes the section that lists other symbols the current symbol uses (the "uses" section).

The 0x8 bit includes the section that lists those symbols that use the current symbol (the "used by" section).

These similar bits apply when you use the `browse-symbol` command, or keys like `<Enter>` or `L` in the `#symbols#` buffer:

The 0x10 bit lets Epsilon include actual source file lines in the `#symbols#` buffer. If off, the buffer will just show their file names and line numbers.

The 0x20 bit includes references to the symbol, not just definitions, when it lists source file lines. That is, it includes the section labeled "is used at".

The 0x40 bit includes the section that lists other symbols the current symbol uses (the "uses" section).

The 0x80 bit includes the section that lists those symbols that use the current symbol (the "used by" section).

Finally, the 0x100 bit applies no matter how the `#symbols#` buffer is built. It lets Epsilon adjust the source line number recorded in the browser database by examining the original source code file. Sometimes the line number recorded in the browser database is off by a line or two, and this option corrects for that.

<code>buf-accessed</code>	System	Buffer-specific	Default: none
---------------------------	--------	-----------------	---------------

Epsilon uses this variable to remember which buffer was accessed most recently. Older buffers have lower values. Each time you switch to a new buffer, Epsilon increments `buf-accessed-clock` and stores it as the new buffer's setting for `buf-accessed`.

<code>buf-accessed-clock</code>	System	Default: none
Epsilon uses this variable to remember which buffer was accessed most recently. See <code>buf-accessed</code> .		
<code>bufed-column-width</code>	Preference	Default: 18
The <code>bufed</code> display leaves this much space for each buffer's name.		
<code>bufed-grouping</code>	Preference	Default: 0
Epsilon can subdivide the list of buffers displayed by the <code>bufed</code> command, and sort each group separately. First it lists buffers with associated files, as well as buffers created by the <code>new-file</code> command. Then it lists buffers without files, typically <code>dired</code> buffers and other special-purpose buffers. Finally (and only if you invoked <code>bufed</code> with a numeric argument), Epsilon lists "system" buffers. The 1 bit in this variable tells Epsilon to separately sort each of these groups according to the current <code>bufed</code> sorting rules, and present the groups in the order shown. Otherwise, Epsilon sorts all groups together, according to the current <code>bufed</code> sorting rules. The 2 bit in this variable, toggled by <code>bufed</code> 's <code>m</code> subcommand, puts modified buffers at the top of the list.		
<code>bufed-show-absolute-path</code>	Preference	Default: 0
Set this variable to 1 if you want the <code>bufed</code> window to display the absolute pathname of each file. A setting of 0 makes Epsilon display a relative path. Under Windows, a setting of 2 makes Epsilon display the 8.3 form of the file name.		
The value 3 makes Epsilon always display the base name of the file first, then the directory portion inside square brackets.		
The value 4 combines options 0 and 3. For files that can be shown with a relative path (those within the current directory set by the <code>cd</code> command, or its subdirectories), <code>bufed</code> shows a relative pathname, followed by the name of the current directory inside square brackets. For other files, <code>bufed</code> shows the full path.		
<code>bufed-width</code>	Preference	Default: 50
This variable contains the width of the pop-up window that the <code>bufed</code> command creates. (Epsilon for Windows doesn't use this variable; instead drag a dialog's border to resize it.)		
<code>buffer-color-scheme</code>	System Buffer-specific	Default: 0
If the buffer-specific variable <code>buffer_color_scheme</code> is non-zero in a buffer, Epsilon uses its value in place of the <code>selected_color_scheme</code> variable when displaying that buffer. It takes precedence over the similar <code>window_color_scheme</code> variable.		
<code>buffer-not-saveable</code>	System Buffer-specific	Default: 0
Some buffers may have an associated file name but should never be automatically saved to that file name or counted as an unsaved file. Bits in this buffer-specific variable control this behavior. The 2 bit (<code>NOSAVE_SESSION</code>) says not to record the file in a session file before exiting Epsilon. The 4 bit (<code>NOSAVE_AUTOSAVE</code>) says not to autosave the file. The 8 bit (<code>NOSAVE_SAVEALL</code>) says that <code>save-all-buffers</code> should skip the file. The 16 bit (<code>NOSAVE_WARN</code>) says not to warn when exiting Epsilon if the file is marked as unsaved. The 32 bit (<code>NOSAVE_BUFED</code>) tells <code>bufed</code> not to list the file, but this only applies if its buffer name starts with <code>-</code> to mark it as a scratch buffer. The 1 bit (<code>NOSAVE_ALL</code>) is the same as specifying all the bits above.		
Also see the <code>force-save-as</code> variable.		

bufname	System	Default: "startup"
---------	--------	--------------------

This variable contains the name of the current buffer. Setting it in an EEL program switches to a different buffer. If the indicated buffer does not exist, nothing happens. Use this method of switching buffers only to temporarily switch to a new buffer; use the `to_buffer()` or `to_buffer_num()` subroutines to change the buffer a window will display.

EEL code must use assignment, not `strcpy()` or similar, to set it: `bufname = "name"`. It's not necessary to preserve the source character array; Epsilon will copy the value.

bufnum	System	Default: none
--------	--------	---------------

This variable contains the number of the current buffer. Setting it in an EEL program switches to a different buffer. If the indicated buffer does not exist, nothing happens. Use this method of switching buffers only to temporarily switch to a new buffer; use the `to_buffer()` or `to_buffer_num()` subroutines to change the buffer a window will display.

build-first	Window-specific	Default: 0
-------------	-----------------	------------

Epsilon normally displays each window line by line, omitting lines that have not changed. When a command has moved point out of the window, Epsilon must reposition the display point (the buffer position at which to start displaying text) to return point to the window. However, Epsilon sometimes does not know that repositioning is required until it has displayed the entire window. When it discovers that point is not in the window, Epsilon moves the display point to a new position and immediately displays the window again. Certain commands which would often cause this annoying behavior set the `build-first` variable nonzero to prevent it.

byte-extension		Default: ".b"
----------------	--	---------------

This variable holds the correct extension of bytecode files in this version of Epsilon.

c-access-spec-offset	Preference Buffer-specific	Default: 0
----------------------	----------------------------	------------

In C mode, Epsilon offsets the indentation of an access specifier (`public:`, `private:`, or `protected:`) by the value of this variable.

c-align-break-with-case	Preference Buffer-specific	Default: 0
-------------------------	----------------------------	------------

If this variable is nonzero, C mode aligns each break statement the same as the previous case statement.

c-align-contin-lines	Preference Buffer-specific	Default: 48
----------------------	----------------------------	-------------

By default, the C indenter tries to align continuation lines under parentheses and other syntactic items on prior lines. If Epsilon can't find anything on prior lines to align with, or if aligning the continuation line would make it start past column `c-align-contin-lines`, Epsilon uses a fixed indentation: two levels more than the original line, plus the value of the variable `c-contin-offset` (normally zero).

Set this variable to zero if you don't want Epsilon to ever try to align continuation lines under syntactic features in previous lines. If zero, Epsilon indents continuation lines by one level (normally one tab stop), plus the value of the variable `c-contin-offset` (which may be negative).

Set this variable to -1 if you want it to align under syntactic elements no matter how far to the right that would go. (This value acts like an infinitely large value.)

`c-align-contin-max-offset` Preference Buffer-specific Default: 0

Set this variable nonzero if you want Epsilon to reduce a continuation line's indentation based on the total width of that line. Set it to -1 to use the current window's width, or a positive value to limit line width to that many columns. (Compare `c-align-contin-lines`, which reduces a continuation line's indentation based on a maximum indentation width, not the total width of a line.)

When a continuation line is wider than the specified number of columns, Epsilon uses the value in the `c-align-contin-max-offset` variable to determine its indentation.

`c-align-contin-max-width` Preference Buffer-specific Default: 0

When a continuation line is longer than the limit set by the `c-align-contin-max-offset` variable, C mode uses this variable to choose a new indentation for the long line.

If greater than zero, Epsilon indents by that amount past the base line (similar to how `c-contin-offset` works). If zero, Epsilon right-aligns the wide line to the column specified by the `c-align-contin-max-width` variable. If negative, it right-aligns but with that amount of extra space.

`c-align-extra-space` Preference Buffer-specific Default: 2

When C mode indents a continuation line, it tries to line up text under previous syntactic constructs. For instance, it may position text just after a `(` character on the previous line. Sometimes (commonly with continued `if` statements) this causes the continuation line to be indented to the same column as following lines. If Epsilon thinks this will happen, it adds the additional indentation specified by this variable to the continuation line.

`c-align-inherit` Preference Buffer-specific Default: 4

When Epsilon indents a continued declaration line that uses C++ inheritance syntax, it indents this much more than the original line. A value of -1 treats it the same as other continuation lines.

`c-align-open-paren` Preference Buffer-specific Default: 0

This variable controls the way Epsilon indents lines that contain only a `(` left parenthesis character. If one, Epsilon aligns it with the start of the current statement. If two, it adjusts the alignment based on the previous line. If zero, it uses extra indentation like other types of continuation lines.

`c-align-selectors` Preference Buffer-specific Default: 48

When Epsilon indents a multi-line Objective-C message expression or definition using selectors (argument labels), it right-aligns the selector on each line with the first selector, as long as they don't go past the column specified by this variable. If they do, or if this variable is zero, it doesn't try to align selectors in this manner.

`c-auto-fill-mode` Preference Buffer-specific Default: 1

Epsilon can break long C/C++/Java/EEL comments as you type them, using a variation of auto-fill mode. Set this variable to 0 to disable this feature. Set it to 2 to let Epsilon break all comments. The default value of 1 tells Epsilon not to break comments that follow non-comment text on the same line, but permit Epsilon to break comments on other lines.

`c-auto-show-delim-chars` Default: "[()]"

This variable holds the set of delimiter characters that should trigger Epsilon's auto-show-delimiters feature in C mode. Epsilon will search for and highlight the match of each delimiter.

`c-biggest-declarator` Preference Buffer-specific Default: 20000

When C mode displays the name of the current function in the mode line, it assumes the count of characters between the function name and the opening brace character that starts its definition is no greater than this value.

`c-block-macro-close-pat` Preference Default: (omitted)

When permitted by the `use-c-macro-rules` variable, Epsilon treats a macro matching this pattern like a close brace character, decreasing the indent level on the following lines.

`c-block-macro-inner-pat` Preference Default: (omitted)

When permitted by the `use-c-macro-rules` variable, Epsilon treats a macro matching this pattern as if it were a normal statement ending in a semicolon, even if the line doesn't end in one.

`c-block-macro-open-pat` Preference Default: (omitted)

When permitted by the `use-c-macro-rules` variable, Epsilon treats a macro matching this pattern like an open brace character, increasing the indent level on the following lines.

`c-brace-offset` Preference Buffer-specific Default: 0

In C mode, Epsilon offsets the indentation of a left brace on its own line by the value of this variable. The `closeback` variable also helps to control this placement.

`c-case-offset` Preference Buffer-specific Default: 0

In C mode, Epsilon offsets the indentation of a case statement by the value of this variable.

`c-color-preproc-if-pattern` Preference Default: (omitted)

In C mode, Epsilon colors blocks that have been commented out with `#if 0` using the same color it uses for comments. It does this for any line starting with `#` that matches the regular expression pattern in this variable.

`c-contain-offset` Preference Buffer-specific Default: 0

In C mode, Epsilon offsets its usual indentation of continuation lines by the value of this variable. The variable only affects lines that Epsilon can't line up under the text of previous lines.

`c-delete-trailing-spaces` Preference Buffer-specific Default: 0

If this variable is nonzero, Epsilon will delete any spaces or tabs at the end of each line just before saving a C mode file.

<code>c-extra-keywords</code>	System	Buffer-specific	Default: 3
-------------------------------	--------	-----------------	------------

C mode automatically sets the buffer-specific `c-extra-keywords` variable based on file name extensions, to indicate which identifiers are considered keywords in the current buffer. The value 1 tells Epsilon to recognize C++ keywords when code coloring. The value 2 tells Epsilon to recognize EEL keywords. The values 4 and 8 indicate Java and IDL keywords, respectively. Epsilon always recognizes those keywords common to C, C++, Java, and EEL.

<code>c-fill-column</code>	Preference	Buffer-specific	Default: 72
----------------------------	------------	-----------------	-------------

This variable sets the default fill column for filling comments in C/C++/Java buffers. If positive, Epsilon uses it to initialize the fill column whenever a buffer enters C mode. (Otherwise Epsilon uses the default value of the `margin-right` variable.)

<code>c-indent</code>	Preference	Buffer-specific	Default: 0
-----------------------	------------	-----------------	------------

C mode indents each additional level of nesting by this many columns. If the variable is less than or equal to zero, Epsilon uses the value of `tab-size` instead. Set this variable if you want Epsilon to use one number for displaying tab characters, and a different number for indenting C code. (Epsilon will indent using a combination of spaces and tabs, as necessary.)

<code>c-indent-after-extern-c</code>	Preference	Buffer-specific	Default: 0
--------------------------------------	------------	-----------------	------------

If zero, a block that starts with `extern "C"` receives no additional indentation.

<code>c-indent-after-namespace</code>	Preference	Buffer-specific	Default: 0
---------------------------------------	------------	-----------------	------------

If zero, a block that starts with a `namespace` declaration receives no additional indentation.

<code>c-label-indent</code>	Preference	Buffer-specific	Default: 0
-----------------------------	------------	-----------------	------------

This variable provides the indentation of lines starting with labels in C mode. Normally, Epsilon moves labels to the left margin.

<code>c-look-back</code>	Preference		Default: 100000
--------------------------	------------	--	-----------------

When C mode tries to determine the correct indentation of a line, it looks back in the buffer at previous lines. To prevent long delays, Epsilon gives up if it finds itself looking back more than this many characters, and uses its best indentation guess so far.

<code>c-param-decl</code>	Preference	Buffer-specific	Default: 0
---------------------------	------------	-----------------	------------

Epsilon indents pre-ANSI K&R-style parameter declarations by the number of characters specified by this variable.

<code>c-prototype-macro-names</code>	Preference		Default: [omitted]
--------------------------------------	------------	--	--------------------

When C mode scans for function definitions, it skips over the standard macros `_ANSI_ARGS`, `__P`, and `PROTO` that can appear within their signatures. You can expand this list of macro names by modifying this variable.

`c-quick-scan` Preference Buffer-specific Default: 1

When `display-definition` scans the buffer in C mode to display the name of the current function, it assumes that a right brace character `}` found at the left margin marks the end of a top-level block (such as a function, class, or structure definition). If you must edit code formatted with the `}` of inner blocks also on the left margin, set this variable zero. (On very large files, this can make updating the function name shown on the mode line much slower.)

`c-reindent-previous-line` Preference Buffer-specific Default: 1

This variable controls whether Epsilon reindents the previous line when you press `<Enter>` in C mode.

`c-spell-options` Preference Default: 0

This variable controls the Spell minor mode in C mode. Use the `spell-mode` command to set it to 1, and Epsilon will highlight misspelled words in C strings and comments. See the `default-spell-options` variable for the other bits you can set to customize spell checking in C mode.

`c-tab-always-indent` Preference Buffer-specific Default: 0

By default, if you press `<Tab>` when point is not in the current line's indentation, C mode inserts a tab character instead of recomputing the current line's indentation. If this variable is nonzero, the `<Tab>` key will reindent the current line, regardless of your position on the line. If you press the key again, it will insert an additional tab.

`c-tab-override` Preference Default: -1

If you want the width of a tab character in C mode buffers to be different than in other buffers, set this variable to the desired value. C mode will change the buffer's tab size to the specified number of columns. Setting this variable doesn't change existing buffers; set the `tab-size` variable for that.

`c-tagging-class` System Default: ""

Epsilon uses this variable while tagging C++/Java files to record the name of the current class.

`c-top-braces` Preference Buffer-specific Default: 0

Epsilon indents the braces of the top-level block of a function by the number of characters specified by this variable. By default, Epsilon puts such braces at the left margin.

`c-top-contin` Preference Buffer-specific Default: 3

Epsilon indents continuation lines outside of any function body by the number of characters specified by this variable, whenever it cannot find any text on previous lines to align the continuation line beneath.

<code>c-top-struct</code>	Preference Buffer-specific	Default: 8
When the definition of a top-level structure, union, or class appears over several lines, Epsilon indents the later lines by the number of characters specified in this variable, rather than the value of <code>c-top-contin</code> .		
<code>c-top-template</code>	Preference Buffer-specific	Default: 0
Successive lines of a definition that uses a C++ template will be indented to the column specified by this variable, when no other indenting rules apply.		
<code>call-on-modify</code>	Buffer-specific	Default: 0
If the buffer-specific <code>call-on-modify</code> variable has a nonzero value in a particular buffer, whenever any primitive tries to modify that buffer, Epsilon calls the EEL subroutine <code>on_modify()</code> first.		
<code>can-get-process-directory</code>		Default: varies
Epsilon sets this variable nonzero to indicate that it is able to retrieve current directory information from the concurrent process. Unix versions of Epsilon will set this variable nonzero only after the process has started and its first prompt has appeared.		
<code>capture-output</code>	Preference	Default: 0
If nonzero, Epsilon makes a transcript of console input and output when it runs another program via the push command. Epsilon puts the transcript in a buffer named “process”.		
<code>case-fold</code>	Preference Buffer-specific	Default: 1
If nonzero, Epsilon considers upper case and lower case the same when searching, so a search string of “Word” would match “word” and “WORD” as well. This variable sets the default for a search in each buffer, but when searching you can change case-folding status for that particular search by pressing Ctrl-c.		
<code>catch-mouse</code>	Preference	Default: varies
If nonzero, Epsilon queues up mouse events. If zero, Epsilon ignores the mouse. Right now only the Windows Console version of Epsilon pays attention to this variable. See the <code>mouse-mask</code> variable to disable mouse support in other versions.		
<code>clear-process-buffer</code>	Preference	Default: 0x20
If 1, the commands push, start-process, compile-buffer, and make will each begin by emptying the process buffer. If 0, the commands append to whatever text is already in the process buffer.		
Other bit values instruct individual commands to empty the process buffer: 0x2 affects push, 0x4 affects start-process, 0x8 affects compile-buffer, and 0x10 affects make, while 0x1 overrides these and forces all these commands to clear. You can add these bit values together.		
If Epsilon isn’t set to use the buffer “process” for one of these commands, then that command will ignore this variable. (For instance, <code>compile-in-separate-buffer</code> normally tells Epsilon to use a separate buffer for compilations; this variable won’t affect that.)		
The 0x20 bit makes Epsilon clear ssh buffers when it reuses them.		

`clipboard-access` Preference Default: 1

If this variable is non-zero, all commands that put text on the kill ring will also try to copy the text to the MS-Windows or X11 clipboard. Similarly, the yank command will retrieve any new text from the clipboard before retrieving text from Epsilon's kill ring if this variable is nonzero.

If you're running Epsilon for Unix as a console program, Epsilon ignores the clipboard, just as if this variable were zero.

During a keyboard macro Epsilon also ignores the clipboard contents. Use the `insert-clipboard` or `copy-to-clipboard` commands if you want to access the clipboard from a keyboard macro. Or set `clipboard-access` to 2, forcing Epsilon to use the clipboard even in a keyboard macro.

`clipboard-convert-mac-lines` Preference Default: 2

When Epsilon inserts text from the clipboard into a buffer not in binary mode, it can replace all Return characters (the traditional Mac line-ending character) with Newline characters (Epsilon's line-ending character). Set this variable to 0 to force this behavior off or 1 to force it on. The default setting 2 enables it only on macOS systems.

`clipboard-convert-unicode` Preference Default: 4

Bits in this variable control how Epsilon converts text when accessing the clipboard.

When the 1 bit is set, yanking text from the clipboard converts certain Unicode characters with codes outside the 0–255 range to corresponding characters that have codes within that range. For instance, Unicode EM SPACE U+2003 becomes a space character, and Unicode HYPHEN U+2011 becomes a – character. This conversion turns some characters into ASCII characters in the 32–127 ASCII range, and others into characters in the Extended ASCII range of 128–159 using the Windows-1252 character set.

The 2 bit enables the opposite conversion when putting text onto the clipboard, but only for characters in the range 128–159.

The 4 bit, which is the only conversion enabled by default, only turns Unicode characters into characters in the 32–127 range, and does not convert others.

Epsilon skips such conversions when yanking and pasting from Unicode buffers (typically, those that already contain characters outside the 0–255 range), and whenever not all the characters can be converted this way. It only does these conversions if all characters outside the 0–255 range in the pasted text can be converted to be within that range.

`clipboard-format` Preference Default: 0

By default, when the Windows Console version of Epsilon puts characters on the MS-Windows clipboard and it can't put them there in Unicode format, it lets Windows translate the characters from the DOS/OEM character set to Windows ANSI. Epsilon needs to do this so that national characters display correctly. When Epsilon retrieves characters from the clipboard, it has Windows perform the reverse translation.

But each character set contains some characters that the other does not, so that copying characters in one direction and then back can change the characters. Instead, you can tell Epsilon to copy characters without translating them. Then copying back and forth will never change the characters, but Epsilon for DOS and Windows won't display the same symbols for any character except the original ASCII printable characters (32 to 127).

Setting this variable to 7 makes Epsilon tell Windows that all text in Epsilon is in the OEM character set, and Windows must translate between DOS/OEM and Windows ANSI. Setting the variable to 1 makes Epsilon tell Windows that all text in Epsilon uses the Windows ANSI character set, so no translating is necessary. The default value of zero makes Epsilon translate only when appropriate. (Epsilon uses the value of this variable as the “clipboard format” to ask Windows for; you can see the raw clipboard data Windows uses by setting the variable to other values, if you like.)

Epsilon for Unix uses this variable to decide which of two X11 system clipboards it should use. A zero value makes Epsilon use the “clipboard selection” used by newer software, a value of one makes it use the “primary selection”.

`closeback` Preference Buffer-specific Default: 1

If nonzero, C mode aligns a right brace character that ends a block with the line containing the matching left brace character. If zero, C mode aligns the right brace character with the first statement inside the block. Additionally, the value 2 makes Epsilon indent lines starting with a right parenthesis so they match the start of the line with its matching left parenthesis, not the left parenthesis itself.

`cmd-len` Default: 0

This variable counts the number of keys in the current command. Epsilon resets it to zero each time it goes through the main loop. It doesn’t count mouse keys or other events that appear as keys.

`cmd-line-session-file` System Default: none

If you use the `-p` flag to provide the name of a particular session file, Epsilon puts the name in this variable.

`coldfusion-empty-elements` Default: " | abort | applet | ... (omitted) | "

HTML mode treats ColdFusion elements named in this list as never using an end tag; thus, they will not begin a region with additional indentation. The “cf” beginning each ColdFusion element name is omitted in this list, so if the list contains “abort”, Epsilon will assume that a start tag named “cfabort” doesn’t have an end tag. Each name must be surrounded by ‘|’ characters.

`color-html-look-back` Preference Default: 500,000

When Epsilon begins coloring HTML in the middle of a buffer, it has to determine whether it’s inside a script by searching back. This can be slow in very large HTML files, so Epsilon limits its search by assuming that a script can be no longer than this many characters.

`color-look-back` Preference Default: 0

When Epsilon begins coloring in the middle of a buffer, it has to determine whether it’s inside a comment by searching back for comment characters. If `color-look-back` is greater than zero, Epsilon only looks back over that many characters for a block comment delimiter like `/*` or `*/` before giving up and concluding that the original text is not inside a comment. If you edit extremely large C files with few block comments, you can speed up Epsilon by setting this variable. Any block comments larger than this value may not be colored correctly. A value of zero (the default) lets Epsilon search as far as it needs to, and correctly colors comments of any size.

<code>color-names</code>	System	Default: " black blue ... "
--------------------------	--------	------------------------------

Epsilon recognizes various color names in command files. It stores the names in this variable.

<code>color-whole-buffer</code>	Preference	Default: 0
---------------------------------	------------	------------

Normally Epsilon colors buffers as needed. You can set Epsilon to instead color the entire buffer the first time it's displayed. Set this variable to the size of the largest buffer you want Epsilon to entirely color at once.

<code>coloring-flags</code>	System	Buffer-specific	Default: 0
-----------------------------	--------	-----------------	------------

Epsilon's syntax highlighting functions use this variable to record various types of status. Bits in the variable are specified by macros in `colcode.h`.

Epsilon uses some bits independently of any particular language mode. `COLOR_DO_COLORING` indicates that Epsilon should perform coloring. `COLOR_IN_PROGRESS` means Epsilon is in the middle of coloring; Epsilon uses this bit to detect when a coloring function has aborted due to a programming error; it then disables coloring for that buffer. `COLOR_MINIMAL` records whether minimal coloring (an option in C/C++/Java mode) is in use for that buffer; Epsilon uses it to notice when this setting has changed.

The remaining bits are set by individual language modes. `COLOR_INVALIDATE_FORWARD` indicates that after the user modifies a buffer, any syntax highlighting information after the modified region should be discarded. `COLOR_INVALIDATE_BACKWARD` indicates that syntax highlighting information before the modified region should be discarded. (Without these bits, Epsilon only discards syntax highlighting information that's very close to the modified part of the buffer.)

`COLOR_INVALIDATE_RESETS` tells Epsilon that whenever it invalidates syntax highlighting in a region, it should also set the color of all text in that region to the default of -1.

`COLOR_RETAIN_NARROWING` indicates that coloring should respect any narrowing in effect (instead of looking outside the narrowed area to parse the buffer in its entirety).

`COLOR_IGNORE_INDENT` says that a simple change of indentation shouldn't cause any recoloring. Languages with no column-related highlighting rules may set this for better performance.

<code>column-in-window</code>		Default: none
-------------------------------	--	---------------

On each screen refresh, Epsilon sets this variable to the column of point within the current window, counting from zero. If you switch windows or move point, Epsilon will not update this variable until the next refresh.

<code>comment-begin</code>	Buffer-specific	Default: ";"
----------------------------	-----------------	--------------

When Epsilon creates a comment, it inserts the contents of the buffer-specific variables `comment-begin` and `comment-end` around the new comment.

<code>comment-column</code>	Buffer-specific	Default: 40
-----------------------------	-----------------	-------------

Epsilon creates and indents comments so they begin at this column, if possible.

`comment-end` Buffer-specific Default: none

When Epsilon creates a comment, it inserts the contents of the buffer-specific variables `comment-begin` and `comment-end` around the new comment.

`comment-pattern` Buffer-specific Default: `" ; .* $"`

The comment commands look for comments using regular expression patterns contained in the buffer-specific variables `comment-pattern` (which should match the whole comment) and `comment-start` (which should match the sequence that begins a comment, like `'/*'`).

`comment-region-options` Preference Default: 0

Bits in this variable modify how the `comment-region` command behaves. The 1 bit makes comments indented: instead of adding the commenting sequence (such as `//`) at the start of each line, Epsilon will add it after each line's existing indentation (but if the lines in the region have varying indentation, the commenting sequences will be aligned based on the line with the least). The 2 bit ensures that removing commenting removes such indented comments too. The 4 bit makes commenting and uncommenting skip blank lines.

A mode may define a mode-specific version of this variable, constructed using its mode name, such as `python-comment-region-options`. If such a variable exists for the buffer's mode, Epsilon will use its value instead.

`comment-repeat-indentation-lines` Preference Buffer-specific Default: 2

Modes that provide language-sensitive indenting, such as C mode (for C, C++, Java, and EEL) and Perl mode, typically indent single-line comments (such as C++'s `//` comments) to the same indentation level as code. Sometimes in the midst of a block of indented code, you may wish to write a series of comment lines with some different indentation.

When Epsilon notices that the 2 previous lines are comment lines, its auto-indenter decides that a blank line that follows should be indented like them, and not as if the line will contain code. Set this variable to change the number of comment lines Epsilon checks for. Set it to zero to make Epsilon always indent blank lines based on language syntax rules.

`comment-start` Buffer-specific Default: `" ; [\t] *"`

The comment commands look for comments using regular expression patterns contained in the buffer-specific variables `comment-pattern` (which should match the whole comment) and `comment-start` (which should match the sequence that begins a comment, like `'/*'`).

`common-open-curdir` System Default: none

In Windows, Epsilon uses this variable to maintain the current directory used in the Common File Dialog that appears when you use the File/Open, File/Save As, or similar menu or toolbar commands.

EEL code must use assignment, not `strcpy()` or similar, to set it: `common_open_curdir = "name"`. It's not necessary to preserve the source character array; Epsilon will copy the value.

`common-open-use-directory` Preference Default: 0

Bits in this variable modify how Epsilon for Windows treats the current directory when it uses a common file dialog (as with the File/Open menu command and similar).

By default, the dialog starts in Epsilon's current directory (which is used to display relative pathnames, among other things, and is set with the `cd` command), the first time you display the common dialog. It is then maintained separately. Navigating in the common file dialog doesn't change Epsilon's current directory. And changing Epsilon's current directory doesn't change the starting directory of the next common file dialog. Epsilon saves the common file dialog's directory from session to session.

Set the variable to 3 if you want this dialog to follow and modify Epsilon's global current directory, as set with the `cd` command. Set it to 1 if you want the dialog to start in Epsilon's current directory, but not to alter it. Set it to 2 if you want the dialog to maintain its own current directory and set Epsilon's current directory to match whenever you select a file.

Set it to 4 if you want the dialog to start in the current buffer's associated directory, if it has one, or the same directory as last time if not. Set it to 8 if want the dialog to start in the current buffer's associated directory if it has one, or Epsilon's global current directory if not. With these last two options, Epsilon's global current directory won't be modified by the dialog.

`compare-files-ignore-spaces` Preference Default: 2

This variable determines how the `compare-files` command compares files that differ only by spaces, tabs, and newlines. The value 0 compares files exactly. The value 1 treats any run of one or more spaces and tabs as a match for any other run of one or more spaces and tabs. The value 2 treats any run of one or more spaces, tabs, or newlines as a match for any other run of one or more spaces, tabs, or newlines. (This variable is ignored when `compare-files-method` is zero.)

`compare-files-method` Preference Default: 2

This variable determines how the `compare-files` command compares files. The value 0 makes it consider files the same if they have the same last-modified timestamp and the same size in bytes on disk. The value 1 makes it consider files the same if their contents are byte-for-byte identical. The value 2 makes it consider files the same if their contents match after the default conversions for line translation and Unicode decoding are performed. (For the latter two options, also see the `compare-files-ignore-spaces` variable.)

`compare-to-prior-version-style` Preference Default: 1

The `compare-to-prior-version` command uses this variable to determine what type of difference listing to generate. A value of 0 produces a listing like that of the `diff` command. A value of 1 produces a listing like `visual-diff`. A value of 2 produces a listing like `merge-diff`.

`compare-windows-ignores-space` Preference Default: 2

This variable says whether the `compare-windows` command should consider any run of one or more whitespace characters in one buffer to match a run of one or more whitespace characters in the other. If 0, it doesn't, and requires all characters to match. If 1, it merges runs of spaces, tabs and newlines. If 2, it merges runs of spaces and tabs only. If 3, it completely ignores differences of only spaces, tabs, and newlines, so `abcdef` matches `abc def`. If 4, it completely ignores differences of only spaces and tabs, but not newlines.

<code>compile-asm-cmd</code>	Preference	Default: <code>m1 "%r"</code>
<p>Epsilon uses the command line contained in the <code>compile-asm-cmd</code> variable to compile Assembly files; those files ending with a <code>.asm</code> extension. See <code>compile-c-cmd</code> for details on this variable's format.</p>		
<code>compile-buffer-cmd</code>	Buffer-specific	Default: none
<p>The <code>compile-buffer</code> command retrieves the command to compile the current buffer from this buffer-specific variable. For C, C++, and EEL files this variable normally just points to the <code>compile-c-cmd</code>, <code>compile-cpp-cmd</code>, or <code>compile-eel-cmd</code> variables, respectively. To make all files with one of the above extensions use a different compile command, set one of these other variables. To make only the current buffer begin to use a different compile command, set this variable.</p> <p>See <code>compile-c-cmd</code> for details on this variable's format.</p>		
<code>compile-c-cmd</code>	Preference	Default: <code>c1 "%r"</code>
<p>Epsilon uses the command line contained in the <code>compile-c-cmd</code> variable to compile C files; those files ending with a <code>.c</code> extension. (Epsilon for Unix uses the <code>compile-c-cmd-unix</code> variable instead.)</p> <p>The command line works as a file name template, so you can substitute parts of the file name into the command line. The sequence <code>%p</code> substitutes the path part of the file name, the sequence <code>%b</code> substitutes the base name (without path or extension), the sequence <code>%e</code> substitutes the extension (including the <code>."</code>), the sequence <code>%f</code> substitutes the full name as an absolute pathname, and the sequence <code>%r</code> substitutes a pathname relative to the current directory. The sequence <code>%x</code> substitutes the full pathname of the directory containing the Epsilon executable. The sequence <code>%X</code> substitutes the same pathname to the Epsilon executable, but converts all Windows long file names to their equivalent short name aliases.</p>		
<code>compile-c-cmd-unix</code>	Preference	Default: <code>cc "%r"</code>
<p>Epsilon for Unix uses the command line contained in this variable to compile C files; those that end with a <code>.c</code> extension. See <code>compile-c-cmd</code> for details on this variable's format, or for the equivalent variable in non-Unix versions.</p>		
<code>compile-cpp-cmd</code>	Preference	Default: <code>c1 "%r"</code>
<p>Epsilon uses the command line contained in the <code>compile-cpp-cmd</code> variable to compile C++ files; those files ending with a <code>.cpp</code> or <code>.cxx</code> extension. See <code>compile-c-cmd</code> for details on this variable's format. (Epsilon for Unix uses the <code>compile-cpp-cmd-unix</code> variable instead.)</p>		
<code>compile-cpp-cmd-unix</code>	Preference	Default: <code>cc "%r"</code>
<p>Epsilon for Unix uses the command line contained in this variable to compile C files; those that end with a <code>.cpp</code> or <code>.cxx</code> extension. See <code>compile-c-cmd</code> for details on this variable's format. See <code>compile-cpp-cmd</code> for the equivalent variable in non-Unix versions.</p>		
<code>compile-csharp-cmd</code>	Preference	Default: <code>csc "%r"</code>
<p>Epsilon uses the command line contained in the <code>compile-csharp-cmd</code> variable to compile C# files; those files ending with a <code>.cs</code> extension. See <code>compile-c-cmd</code> for details on this variable's format.</p>		

<code>compile-eel-cmd</code>	Preference	Default: <code>%Xeel "%r"</code>
------------------------------	------------	----------------------------------

Epsilon uses the command line contained in the `compile-eel-cmd` variable to compile EEL files; those files ending with a `.e` extension. After an EEL file has been successfully compiled, Epsilon will automatically load it. Epsilon for Windows or Unix generally uses its built-in EEL compiler instead of this variable; see `compile-eel-dll-flags` to set its flags. See `compile-c-cmd` for details on this variable's format.

<code>compile-eel-dll-flags</code>	Preference	Default: <code>"-n -q"</code>
------------------------------------	------------	-------------------------------

When Epsilon compiles EEL code using its internal EEL compiler, it looks in this variable for EEL command line flags.

<code>compile-gams-cmd</code>	Preference	Default: <code>gams "%r"</code>
-------------------------------	------------	---------------------------------

Epsilon uses the command line contained in the `compile-gams-cmd` variable to compile GAMS files; those files ending with a `.gms` extension (or others). See `compile-c-cmd` for details on this variable's format.

<code>compile-html-cmd</code>	Preference	Default:
-------------------------------	------------	----------

The `compile-buffer` command uses the command line contained in the `compile-html-cmd` variable to "compile" an HTML file. In the case of HTML files, this means to run a user-specified command that could perform syntax-checking or some similar function. (To display the file in a browser, use the command `quick-dired-command` by pressing `Alt-o` v.) See `compile-c-cmd` for details on this variable's format.

<code>compile-idl-cmd</code>	Preference	Default: <code>midl "%r"</code>
------------------------------	------------	---------------------------------

Epsilon uses the command line contained in the `compile-idl-cmd` variable to compile IDL files; those files ending with an `.idl` or `.acf` extension. Epsilon normally edits such files using C mode. See `compile-c-cmd` for details on this variable's format.

<code>compile-in-separate-buffer</code>	Preference	Default: 1
---	------------	------------

In some environments, the `compile-buffer` and `make` commands can do their work in a separate compilation buffer. This is the most reliable way for them to work. Set this variable to zero to force them to share an existing process buffer.

By default, when Epsilon compiles in a separate buffer, it doesn't display the buffer if the compilation finishes without errors. Set this variable to 2 if you want Epsilon to display the buffer as soon as the compilation starts.

<code>compile-java-cmd</code>	Preference	Default: <code>javac "%r"</code>
-------------------------------	------------	----------------------------------

Epsilon uses the command line contained in the `compile-java-cmd` variable to compile Java files; those files ending with a `.java` extension. See `compile-c-cmd` for details on this variable's format.

<code>compile-latex-cmd</code>	Preference	Default: (omitted)
--------------------------------	------------	--------------------

Epsilon uses the command line contained in the `compile-latex-cmd` variable to compile LaTeX files; those files ending with a `.ltx`, `.sty`, or (in most cases) `.tex` or `.cls` extension. See `compile-c-cmd` for details on this variable's format.

`compile-makefile-cmd` Preference Default: `nmake /f "%r"`

The `compile-buffer` command uses the command line contained in the `compile-makefile-cmd` variable to “compile” a makefile (those files ending with a `.mak` extension or named makefile). In the case of makefiles, this means to run `make` on it. See `compile-c-cmd` for details on this variable’s format. See `compile-makefile-cmd-unix` for the Unix equivalent.

`compile-makefile-cmd-unix` Preference Default: `make -f "%r"`

The `compile-buffer` command uses the command line contained in the `compile-makefile-cmd` variable to “compile” a makefile (those files ending with a `.mak` extension or named makefile) under Unix. In the case of makefiles, this means to run `make` on it. See `compile-c-cmd` for details on this variable’s format. See `compile-makefile-cmd` for the non-Unix equivalent.

`compile-perl-cmd` Preference Default: `perl "%r"`

The `compile-buffer` command uses the command line contained in the `compile-perl-cmd` variable to “compile” a Perl file (those files ending with a `.perl` extension or others). In the case of Perl files, this means to execute it. See `compile-c-cmd` for details on this variable’s format.

`compile-php-cmd` Preference Default: `"`

The `compile-buffer` command uses the command line contained in the `compile-php-cmd` variable to “compile” a PHP file. In the case of PHP files, this means to run a user-specified command that could perform syntax-checking or some similar function. (To display the file in a browser, use the command `quick-dired-command` by pressing `Alt-o v`.) See `compile-c-cmd` for details on this variable’s format.

`compile-python-cmd` Preference Default: `python "%r"`

The `compile-buffer` command uses the command line contained in the `compile-python-cmd` variable to “compile” a Python file (those files ending with a `.py` or `.jy` extension or others). In the case of Python files, this means to execute it. See `compile-c-cmd` for details on this variable’s format.

`compile-rmode-cmd` Preference Default: `R <"%r"`

Epsilon uses the command line contained in the `compile-rmode-cmd` variable to compile files in the R language. See `compile-c-cmd` for details on this variable’s format.

`compile-tcl-cmd` Preference Default: `tclsh "%r"`

Epsilon uses the command line contained in the `compile-tcl-cmd` variable to compile Tcl files; those files ending with a `.tcl` or `.ttml` extension. See `compile-c-cmd` for details on this variable’s format.

`compile-tex-cmd` Preference Default: (omitted)

Epsilon uses the command line contained in the `compile-tex-cmd` variable to compile TeX files; those files ending with a `.tex` extension that aren’t LaTeX files (but see the `tex-force-latex` variable). See `compile-c-cmd` for details on this variable’s format.

`compile-typescript-cmd` Preference Default: `tsc "%x"`

Epsilon uses the command line contained in the `compile-typescript-cmd` variable to compile TypeScript files, a variant of JavaScript that uses files ending with a `.ts` extension. See `compile-c-cmd` for details on this variable's format.

`compile-vbasic-cmd` Preference Default: `vbc "%x"`

Epsilon uses the command line contained in the `compile-vbasic-cmd` variable to compile Visual Basic files. See `compile-c-cmd` for details on this variable's format.

`compile-vhdl-cmd` Preference Default: `vhdl "%x"`

The `compile-buffer` command uses the command line contained in the `compile-vhdl-cmd` variable to compile a VHDL file (those files ending with a `.vhdl` extension or others). See `compile-c-cmd` for details on this variable's format.

`compile-xml-cmd` Preference Default:

The `compile-buffer` command uses the command line contained in the `compile-xml-cmd` variable to “compile” an XML file. In the case of XML files, this means to run a user-specified command that could perform syntax-checking or some similar function. See `compile-c-cmd` for details on this variable's format.

`completion-pops-up` Preference Default: 1

If Epsilon cannot add any letters when you ask for completion on a file name or similar item, it will pop up a list of items that match what you've typed so far. To disable automatic pop-ups on completion, set the `completion-pops-up` variable to zero.

`concurrent-compile` Preference Buffer-specific Default: 3

The buffer-specific `concurrent-compile` variable controls how the `compile-buffer` command behaves. If 0, `compile-buffer` always runs the compiler or other program non-concurrently, exiting the concurrent process if it needs to. If 2, the `compile-buffer` command always runs the compiler concurrently, creating a concurrent process if it needs to. If 1, the `compile-buffer` command runs the compiler concurrently if a concurrent process is already running, non-concurrently otherwise. If 3 (the default), `compile-buffer` uses the value of the `concurrent-make` variable instead.

`concurrent-make` Preference Default: 1

The `concurrent-make` variable controls how the `make` command behaves. If 0, the `make` command always runs the compiler or other program non-concurrently, exiting the concurrent process if it needs to. If 2, the `make` command always runs the compiler concurrently, creating a concurrent process if it needs to. If 1 (the default), the `make` command runs the compiler concurrently if a concurrent process is already running, non-concurrently otherwise.

`conf-auto-show-delim-chars` Default: `"[()]"`

This variable holds the set of delimiter characters that should trigger Epsilon's `auto-show-delimiters` feature in Conf file mode. Epsilon will search for and highlight the match of each delimiter.

`console-ansi-font` Preference Default: 0

The Windows Console version of Epsilon normally uses a font with a DOS/OEM character set. It can only display characters from that character set, not arbitrary Unicode characters, even if the underlying font contains them.

Under Windows NT/2000/XP and later Windows versions, you can set this variable nonzero to make Epsilon display Unicode characters if they're in the console's font. This also makes Epsilon for Windows Console use Unicode rules for 8-bit character classifications (such as whether a particular character number in the range 128–255 represents an uppercase letter) instead of the rules for the character set of the current DOS/OEM code page, and modifies how Epsilon uses the Windows clipboard.

`context-help-default-rule` Preference Default: `"=https://www.google.com/search?q="`

This variable defines the context help rule Epsilon uses when no mode-specific rule is available. See page 102 for details on its format.

`context-help-rule-asm-unix` Preference Default: `"`

This variable specifies the context help rule Epsilon uses in Asm mode under Unix. See page 102 for details on its format.

`context-help-rule-asm-windows` Preference Default: (omitted)

This variable specifies the context help rule Epsilon uses in Asm mode under Windows. See page 102 for details on its format.

`context-help-rule-c-unix` Preference Default: `"context_help_man"`

This variable specifies the context help rule Epsilon uses in C mode under Unix. See page 102 for details on its format.

`context-help-rule-c-windows` Preference Default: (omitted)

This variable specifies the context help rule Epsilon uses in C mode under Windows. See page 102 for details on its format.

`context-help-rule-eel` Preference Default: `">epsilon"`

This variable specifies the context help rule Epsilon uses for EEL files. See page 102 for details on its format.

`context-help-rule-gams` Preference Default: `"`

This variable specifies the context help rule Epsilon uses in GAMS mode. See page 102 for details on its format.

`context-help-rule-html` Preference Default: (omitted)

This variable specifies the context help rule Epsilon uses in HTML mode. See page 102 for details on its format.

`context-help-rule-java` Preference Default: (omitted)

This variable specifies the context help rule Epsilon uses in Java mode. See page 102 for details on its format.

`context-help-rule-jscript` Preference Default: (omitted)

This variable specifies the context help rule Epsilon uses in JavaScript mode. See page 102 for details on its format.

`context-help-rule-latex` Preference Default: (omitted)

This variable specifies the context help rule Epsilon uses in LaTeX mode. See page 102 for details on its format.

`context-help-rule-nsis` Preference Default: (omitted)

This variable specifies the context help rule Epsilon uses in NSIS mode. See page 102 for details on its format.

`context-help-rule-objc` Preference Default: (omitted)

This variable specifies the context help rule Epsilon uses in Objective-C mode. See page 102 for details on its format.

`context-help-rule-perl` Preference Default: "+context_help_perldoc"

This variable specifies the context help rule Epsilon uses in Perl mode. See page 102 for details on its format.

`context-help-rule-php` Preference Default: (omitted)

This variable specifies the context help rule Epsilon uses in PHP mode. See page 102 for details on its format.

`context-help-rule-postscript` Preference Default: ""

This variable specifies the context help rule Epsilon uses in PostScript mode. See page 102 for details on its format.

`context-help-rule-python` Preference Default: (omitted)

This variable specifies the context help rule Epsilon uses in Python mode. See page 102 for details on its format.

`context-help-rule-shell` Preference Default: "+context_help_man"

This variable specifies the context help rule Epsilon uses in Shell mode. See page 102 for details on its format.

`context-help-rule-tcl` Preference Default: (omitted)

This variable specifies the context help rule Epsilon uses in Tcl mode. See page 102 for details on its format.

`context-help-rule-tex` Preference Default: ""

This variable specifies the context help rule Epsilon uses in TeX mode. See page 102 for details on its format.

`context-help-rule-vbasic-unix` Preference Default: ""

This variable specifies the context help rule Epsilon uses in Visual Basic mode under Unix. See page 102 for details on its format.

`context-help-rule-vbasic-windows` Preference Default: (omitted)

This variable specifies the context help rule Epsilon uses in Visual Basic mode under Windows. See page 102 for details on its format.

`context-help-rule-vbscript` Preference Default: (omitted)

This variable specifies the context help rule Epsilon uses in VBScript mode. See page 102 for details on its format.

`context-help-rule-vhdl` Preference Default: (omitted)

This variable specifies the context help rule Epsilon uses in VHDL mode. See page 102 for details on its format.

`context-help-rule-xml` Preference Default: None

This variable specifies the context help rule Epsilon uses in XML mode. See page 102 for details on its format.

`context-help-site-search` Preference Default: (omitted)

When a context help rule contains & followed by a website host name, Epsilon constructs a URL from this variable, substituting the host name for its @ character. Then it replaces its * character with the keyword to be looked up, and has a web browser open that URL. It should show a list of matches. See page 102 for more details.

`context-help-site-search-lucky` Preference Default: (omitted)

When a context help rule contains @ followed by a website host name, Epsilon constructs a URL from this variable, substituting the host name for a later @ character in the value. Then it replaces its * character with the keyword to be looked up, and has a web browser open that URL. It should go to the first web page in the list of matches. See page 102 for more details.

`copy-include-file-name-batch` Preference Default: "call %r"

In Batch mode (used for Windows .bat files and similar), the `copy-include-file-name` command uses this variable to construct a line that will call the current file, when inserted into some other batch file. The variable holds a file name template (see page 122) which is used to format the current file's name.

`copy-include-file-name-latex` Preference Default: "\RequirePackage{%b}"

In LaTeX mode, the `copy-include-file-name` command uses this variable to construct a line that will input the current file, when inserted into some other LaTeX file. The variable holds a file name template (see page 122) which is used to format the current file's name.

`copy-include-file-name-options` Preference Default: 0

This file contains bits that alter the behavior of the `copy-include-file-name` command. Set the 1 bit to make the command always use the syntax `#include "file"` in C mode. By default, it uses the syntax `#include <file>` for files it finds on the include path.

`copy-include-file-name-perl` Preference Default: "use %b;"

In Perl mode, the `copy-include-file-name` command uses this variable to construct a line that will import the current file, when inserted into some other Perl file. The variable holds a file name template (see page 122) which is used to format the current file's name.

`copy-include-file-name-shell` Preference Default: "source %r"

In Shell mode, the `copy-include-file-name` command uses this variable to construct a line that will call the current file, when inserted into some other shell script. The variable holds a file name template (see page 122) which is used to format the current file's name.

`copy-include-file-name-tex` Preference Default: "\input{%b}"

In TeX mode, the `copy-include-file-name` command uses this variable to construct a line that will input the current file, when inserted into some other TeX file. The variable holds a file name template (see page 122) which is used to format the current file's name.

`css-auto-show-delim-chars` Default: "[()]"

This variable holds the set of delimiter characters that should trigger Epsilon's auto-show-delimiters feature in CSS mode, or in CSS code embedded in an HTML file. Epsilon will search for and highlight the match of each delimiter.

`css-indent` Preference Buffer-specific Default: 8

CSS mode indents by this many columns for each additional level of nesting.

`current-output-color` System Buffer-specific Default: varies

Process buffer coloring uses this variable internally to interpret ANSI color escape sequences and apply the color. It contains the color class number set by the most recent escape sequence.

<code>cursor-blink-period</code>	Preference	Default: 100
----------------------------------	------------	--------------

This variable controls the rate at which the text cursor blinks. It specifies the period of the on/off cycle in hundredths of a second. It only applies when Epsilon runs as an X11 program in Unix. Set this to -1 to disable blinking.

<code>cursor-shape</code>	System	Default: 98099
---------------------------	--------	----------------

This variable holds the current cursor shape code. Epsilon copies values from `overwrite-cursor`, `normal-cursor`, or one of the other cursor variables, as appropriate, into this variable whenever you switch windows or buffers. Set those variables instead of this one. Epsilon only uses this variable in console environments. See `gui-cursor-shape` for the Windows or Unix equivalent.

<code>cursor-to-column</code>	Window-specific	Default: -1
-------------------------------	-----------------	-------------

The window-specific `cursor-to-column` variable lets you position the cursor in a part of a window where there are no characters. It's normally -1, and the cursor stays on the character after point. If it's non-negative in the current window, Epsilon puts the cursor at the specified column in the window instead. Epsilon resets `cursor-to-column` to -1 whenever the buffer changes, or point moves from where it was when you last set `cursor-to-column`. (Epsilon only checks these conditions when it redisplay the window, so you can safely move point temporarily.)

<code>cygwin-filenames</code>	Preference	Default: 0
-------------------------------	------------	------------

This variable makes Epsilon for Windows recognize file names in the various formats used by Cygwin programs that simulate a Unix-style environment under Windows, when they appear as output in the process buffer and in certain other contexts. Bits in the variable enable Epsilon to translate the different formats. Add the bits together to have Epsilon recognize multiple formats. A setting of 0x60 works well.

In this version, it recognizes the file names when they appear in directory-change messages produced by Gnu Make, or in prompts produced by a Unix-style shell that include the current directory name (so long as they're in the format "dirname>"), in certain file names generated in the `search-man-pages` command, and on Epsilon's command line. Add these bits together to select the rules you want:

The value 0x1 makes Epsilon recognize the format `//c/windows/file` (instead of `c:\windows\file`). This format conflicts with the format for Windows network file names, so servers with one-letter names won't be accessible if you enable this feature. This older format is considered obsolete.

The value 0x2 makes Epsilon recognize the format `/cygdrive/c/windows/file` (instead of `c:\windows\file`).

The value 0x4 makes Epsilon recognize the format `~/file`. Epsilon substitutes the value of the HOME environment variable for the `~`.

The value 0x8 makes Epsilon modify directory names appearing at a prompt that start with `/` by prepending the value of the `cygwin-root` variable.

The value 0x10 is just like the above one for 8, but applies only when Epsilon is parsing Make output messages like "Entering /usr/src/prog".

The value 0x20 makes Epsilon invoke the `cygpath` program on any names not handled by any of the above rules you've enabled.

The value 0x40 makes Epsilon use / instead of \ when <Tab> performs file name completions in the process buffer. It doesn't make completion use Cygwin's other file name format rules, such as using /cygdrive/d instead of d:.

cygwin-root Preference Default: ""

Epsilon for Windows can use this variable to help convert pathnames from Cygwin format. See the cygwin-filenames variable for details.

date-format Preference Default: "%D-%t-%y %h:%N:%E %am"

When the insert-date command inserts the current time and date, or Epsilon displays a time and date in certain other contexts, it formats the date using this variable. It can include any of the following sequences, and Epsilon will substitute the indicated value for that sequence:

%i Epsilon substitutes the current hour in the range 0 to 23.

%h Epsilon substitutes the current hour in the range 1 to 12.

%n Epsilon substitutes the current minute in the range 0 to 59.

%e Epsilon substitutes the current second in the range 0 to 59.

%a Epsilon substitutes "a" before noon, "p" otherwise.

%d Epsilon substitutes the day of the month in the range 1 to 31.

%m Epsilon substitutes the month number in the range 1 to 12.

%f Epsilon substitutes the English month name like January or February.

%t Epsilon substitutes a 3-letter month abbreviation like Jan or Feb.

%y Epsilon substitutes the four digit year.

%Y Epsilon substitutes the two digit year.

%w Epsilon substitutes the name of the day of the week, such as Friday.

%s Epsilon substitutes the name of the day of the week, abbreviated to 3 characters, such as Fri.

%% Epsilon substitutes a literal "%" character.

Using an uppercase letter instead of a lowercase letter for the sequences %i, %h, %n, %e, %d, or %m makes Epsilon always use two digits instead of one. For instance, at 3:07 pm, %h%N produces 3:07, while %H%n produces 03:7.

Using an uppercase letter instead of a lowercase letter for the sequences %a, %f, %w, %s, or %t makes Epsilon insert an uppercase version of the specified month, weekday, or AM/PM indicator.

default-add-final-newline Preference Default: 0

If this variable is nonzero, Epsilon will ensure that each non-empty non-binary file it saves ends with a newline character, by inserting one at the end of the buffer, if necessary, just before saving a file. Before checking this variable, Epsilon looks for a mode-specific variable with a name of the form c-add-final-newline and uses that instead. (Because of line translation, the newline character inserted in the buffer could become a return/linefeed sequence or something else when saved on disk. See set-line-translate for details.)

`default-character-set` Preference Default: 0

Set this variable to 2 if you want Epsilon for Windows to translate character sets by default, in the manner of the `find-oem-file` command. Set it to any other value to disable this behavior.

`default-color-spell-word-pattern` Preference Default: (omitted)

The Spell minor mode uses this regular expression pattern to find misspelled words in those modes that use syntax highlighting and have no spell word pattern of their own. It finds words based on their syntax highlighting, looking for those using a color class name that ends in `-text`, `-comment`, or `-string`. Also see `default-spell-word-pattern`.

`default-delete-trailing-spaces` Preference Default: 0

If this variable is nonzero, Epsilon will delete any spaces or tabs at the end of each line just before saving a file (except for binary files). Before checking this variable, Epsilon looks for a mode-specific variable like `c-delete-trailing-spaces`.

`default-read-encoding` Preference Default: ""

When Epsilon reads a file, no particular encoding was selected for it, and auto-detecting (see the `detect-encodings` variable) doesn't select an encoding, it uses the encoding named by this variable. If the variable doesn't contain the name of a known encoding, Epsilon uses the Raw encoding.

`default-reindent-previous-line` Preference Default: 0

This variable controls whether Epsilon reindents the previous line when you press `(Enter)`, in language modes that include automatic indentation. Before checking this variable, Epsilon looks for a mode-specific variable like `c-reindent-previous-line` or `vbasic-reindent-previous-line`.

`default-spell-options` Preference Default: 0

This variable controls the Spell minor mode in those modes that don't have a dedicated variable for spelling.

Epsilon puts `-spell-options` after the major mode's name and looks for a variable by that name. If none, it uses this variable to see if it should highlight misspelled words. Use the `spell-mode` command to set it to 1, and Epsilon will highlight misspelled words in all such modes.

All `-spell-options` variables contain these bits, which you can set to customize how spell checking works:

Bit	Meaning
0x1	Highlight misspelled words.
0x2	Skip words containing an underscore.
0x4	Skip MixedCaseWords (those with internal capitalization).
0x8	Skip uppercase words (those with no lowercase letters).
0x10	Skip words following a digit, like 14th.
0x20	Skip words before a digit, like gr8.
0x200	Don't remove 's when checking words.
0x1000	Provide faster but less accurate built-in suggestions.
0x2000	Don't copy the case of the original in built-in suggestions.
0x4000	Add globally ignored words to spell helper's list.

When the `spell-helper-program` variable is nonempty so that Epsilon gets its suggestions by running an external program, it ignores the 0x10 and 0x20 bits.

`default-spell-word-pattern` Preference Default: (omitted)

The Spell minor mode uses this regular expression pattern to find misspelled words in those modes that don't use standard syntax highlighting and have no spell word pattern of their own. It doesn't exclude any words based on syntax rules. Also see the `default-color-spell-word-pattern` variable.

`default-state-file-name` System Default: `epsilon-v14.sta`

Epsilon sets this variable to the name of the state file it will look for when it starts. This name contains Epsilon's major version number, so customized state files for different versions of Epsilon can share the same customization directory.

`default-translation-type` Preference Default: 5

When you read an existing file, Epsilon consults this variable to determine what kind of line translation to perform. If 5 (`FILETYPE_AUTO`), Epsilon examines the file's contents and selects one of the following translations, setting the buffer's `translation-type` variable to the selected translation. If this variable is set to any other value, Epsilon uses the specified translation without examining the contents of the file.

A value of 0 (`FILETYPE_BINARY`) makes Epsilon do no line translation, 1 (`FILETYPE_MSDOS`) makes Epsilon strip `<Return>` characters when reading and insert them when writing, 2 (`FILETYPE_UNIX`) makes Epsilon do no line translation, but indicates that the file contains text, 3 (`FILETYPE_MAC`) makes Epsilon replace `<Return>` characters with `<Newline>` characters when reading, and replace `<Newline>` characters with `<Return>` characters when writing.

For remote files (those that use URL syntax), Epsilon uses the `force-remote-translation-type` variable instead. Also see `new-buffer-translation-type` to change the translation rules for newly-created files and buffers.

`default-word` Preference Default: `"<word>+"`

The word commands use a regular expression to define the current notion of a word. While a mode can provide its own regular expression for words, most modes use the regular expression found in this variable in versions of Epsilon for Windows and Unix.

`default-write-encoding` Preference Default: `"UTF-8-No-BOM"`

When Epsilon saves a buffer that has some 16-bit Unicode characters, but no encoding is set for the buffer, it normally prompts for the encoding to use. If it must save a file in a context where prompting wouldn't be appropriate (auto-saving, for example), it falls back to this encoding.

`delete-hacking-tabs` Preference Buffer-specific Default: 0

If 1, when `<Backspace>` deletes a tab, it first turns the tab into the number of spaces necessary to keep the cursor in the same column, then deletes one of the spaces. If 2, when `<Backspace>` deletes a space, it deletes additional spaces and tabs until it reaches the previous tab column (set

by `soft-tab-size` or `tab-size`). The first setting makes `<Backspace>` treat tabs more like spaces; the second makes it treats spaces more like tabs.

You can add these bit settings to get both behaviors. Other bits restrict when deleting a space deletes additional spacing to reach a tab column. The bit 4 does this only within a line's indentation (in the run of whitespace at the start of a line).

The 8 bit deletes multiple spaces only when the deletion starts in a tab column. So if tabs are every 4 columns and a line has 10 spaces, pressing `<Backspace>` three times starting in column 10 will delete one space, one space, then four spaces.

The 0x10 bit deletes multiple spaces only when in the middle of spacing; there must be spaces or tabs both before and after point.

The `<Backspace>` key (or to be precise, the `backward-delete-character` command it runs by default) never deletes multiple spaces due to `delete-hacking-tabs` when given a numeric prefix argument.

Before `backward-delete-character` checks this variable, it looks for a mode-specific one, a variable whose name consists of the current mode name followed by `-backward-delete-character`. If found, it uses that variable's value instead.

`delete-options` Preference Default: 0

Bits in this variable controls optional behavior in deleting and killing commands. Normally, if a command that deletes single characters, like the one on `<Backspace>` or ``, follows a command that puts text on the kill ring, the additional characters will be added to the kill ring. The 1 bit prevents that; the characters will simply be deleted. With a numeric argument, these commands delete many characters at once; in that case they always put their text on the kill ring.

`detect-encodings` Preference Default: 127

Bits in this variable control which Unicode encodings Epsilon tries to auto-detect when it reads a file (when a file doesn't start with a byte order mark that specifies a particular encoding).

A value of 1 makes Epsilon autodetect most valid UTF-8 files. A value of 2 or 4 makes Epsilon autodetect some UTF-16LE and UTF-16BE files, respectively.

A value of 8 makes Epsilon scan the file's first line (treating it as plain ASCII) for a file variable named "coding" that specifies the correct encoding. It then interprets the entire file using the specified encoding. The coding value specified must match a known encoding name or alias, or it will be ignored.

By default, Epsilon tries to autodetect all of these.

`diff-match-characters` Preference Default: 5

When the `visual-diff` command highlights runs of modified characters within each group of modified lines, it ignores short runs of matching characters. This variable specifies the size of the smallest run of matching characters it will recognize.

`diff-match-characters-limit` Preference Default: 50000

The `visual-diff` command highlights runs of modified characters within each group of modified lines. To avoid long delays, it does this only if both runs of modified lines are smaller than this size in characters. If either run contains this many characters or more, `visual-diff` presents that group of lines without character-based highlighting. Set this variable to zero to entirely disable `visual diff`'s highlighting based on individual characters; highlighting will then always be line-based.

<code>diff-match-lines</code>	Preference	Default: 1
When resynchronizing, diff believes it has found another match when <code>diff-match-lines</code> lines in a row match.		
<code>diff-mismatch-lines</code>	Preference	Default: 500
When resynchronizing, diff gives up if it cannot find a match within <code>diff-mismatch-lines</code> lines.		
<code>diff-precise-limit</code>	Preference	Default: 4,000,000
The diff command normally uses an algorithm that finds the minimum set of differences between the lines of two buffers. But this algorithm becomes slow on very large buffers. So if both buffers are larger (in bytes) than this setting, Epsilon uses a different algorithm that doesn't always find the absolute minimum set of differences (and may give up if the buffers are too different, according to <code>diff-mismatch-lines</code>), but is much faster.		
<code>directory-flags</code>		Default: 0
When you specify the <code>-w</code> flag on the command line, Epsilon puts its numeric parameter in this variable.		
<code>dired-24-hour-time</code>	Preference	Default: 2
Set this variable to 1 if you want the dired command in non-Unix versions of Epsilon to display times in 24-hour format. Set it to 0 if you want 12-hour format with AM and PM indicators. The value 2 makes Epsilon for Windows use the system's setting for this.		
<code>dired-buffer-pattern</code>	System Buffer-specific	Default: none
When dired wants to rebuild the file list in the current dired buffer, it looks in this variable for the directory name or file pattern to use. If this variable is null, it uses the name of the dired buffer as the pattern.		
<code>dired-confirmation</code>	Preference	Default: 127
Bits in this variable control when dired's subcommands issue an additional prompt for confirmation. The 1 bit makes dired prompt before deleting a read-only file. The 2 bit makes dired warn before a copy operation overwrites a local file, and the 4 bit does the same for renaming operations. The 8 bit makes dired prompt before deleting a directory hierarchy.		
<code>dired-format</code>	System Buffer-specific	Default: 0
Running dired on a remote directory of files uses this variable to record the format of the directory listing. The variable is zero for local directories in Epsilon's standard format.		
<code>dired-groups-dirs</code>	System	Default: 1
The dired-sort command uses the <code>dired-groups-dirs</code> variable to record whether or not to group subdirectories. If nonzero, all subdirectories appear in a dired listing before any of the files in that directory. If zero, the subdirectories are sorted in with the files, except for the <code>.</code> and <code>..</code> subdirectories, which always appear first regardless of this setting. Use the S key in a dired buffer to set this variable.		

`dired-layout` Preference Default: 1

Most bits in this variable control which file attributes Epsilon for Windows displays in its directory listings for the `dired` command, as follows: 0x1: Read-only, 0x2: Hidden, 0x4: System, 0x8: Archive. Add the values together to include columns for multiple attributes. A value of 0xf displays all attributes. The 0x10 bit makes Epsilon for Windows use the format 2005-12-31 for file dates instead of the usual one.

The 0x20 bit makes directory listings on macOS, Linux and FreeBSD show the full date including year and seconds for all files. (By default Epsilon mimics the format of “ls -l”, showing either a year or hours and minutes, never both and never seconds.)

`dired-live-link-limit` Preference Default: 5,000,000

Dired’s live link feature shows the contents of files in a separate window as you move about in the dired buffer. To prevent long delays, it skips automatically showing files bigger than this many bytes.

`dired-show-dotfiles` Preference Default: 3

The `dired` command can hide files and directories whose names start with a period character. The `-` subcommand toggles whether such items are hidden at all, while bits in this variable select which items will still be shown even when hiding is enabled. The bit 1 means the `.` entry for a directory will always be shown, and the bit 2 means the `..` entry will always be shown. The 4 bit means all other directory names will be shown, even if their names start with a period. The 8 bit means non-directories (including ordinary files) will be shown, even if their names start with a period.

For `ftp://` files, see the `ftp-compatible-dirs` variable. For `scp://` files, note that certain sftp servers may not display hidden files under any circumstances.

`dired-sorts-files` System Default: 'n'

The `dired-sort` command uses the `dired-sorts-files` variable to record how to sort dired buffers. It contains a letter code to indicate the type of sorting: N, E, S, or D to sort by file name, file extension, size, or time and date of modification, respectively, or the value 0 to leave the listing unsorted. An upper case letter code indicates a descending (reverse) sort, a lower case letter code indicates the normal ascending sort. Set this variable using dired’s `S` subcommand.

`discardable-buffer` Buffer-specific Default: 0

Epsilon warns you before exiting if any “valuable” unsaved buffers exist. It considers a buffer valuable if it has a file name associated with it and contains at least one character. An EEL program can set this buffer-specific variable to a nonzero value to indicate that the current buffer doesn’t require any such warning.

`display-c-options` Preference Default: 1

When Epsilon displays the current function’s name on the mode line, it normally includes any class and namespace for the function. Set this variable to zero to omit any namespace name if there’s a class.

`display-column` Preference Window-specific Default: 0

This variable determines how Epsilon displays long lines. If negative, Epsilon displays buffer lines too big to fit on one screen line on multiple screen lines, with a special character to indicate that the line has been wrapped. If `display-column` is 0 or positive, Epsilon only displays the part of a line that fits on the screen. Epsilon also skips over the initial `display-column` columns of each line when displayed. Horizontal scrolling works by adjusting the display column.

`display-definition` Preference Default: 1

In C/C++/Java/Perl buffers, among others, Epsilon can display the name of the current function, subroutine, class, or structure on a buffer's mode line, or in the title bar of Epsilon's window. Set this variable to 2 if you want Epsilon to use the title bar if possible. Versions of Epsilon that can't set the title bar will instead use the mode line. Set this variable to 1 if you want to use the mode line regardless. Or set this variable to 0 to disable this feature. You can modify the `mode-format` variable to position the name within the mode line.

On very large buffers, Epsilon won't update this information while you drag the scroll bar, only when you release it. See the `mouse-update-during-scroll` variable.

`display-func-name` System Default: none

Epsilon uses this variable to help display the name of the current function on the mode line or window title bar. It contains the most recent function name Epsilon found.

`display-scroll-bar` System Window-specific Default: 0

This variable controls whether the current window's right border contains a scroll bar. Set it to zero to turn off the scroll bar, or to any positive number to display the bar. If a window has no right border, or has room for fewer than two lines of text, Epsilon won't display a scroll bar. Although the EEL functions that come with Epsilon don't support clicking on a scroll bar on the left border of a window, Epsilon will display one if the variable is negative. Any positive value produces the usual right-border scroll bar. Run the `toggle-scroll-bar` command instead of setting this internal variable directly.

`double-click-time` Preference Default: 40

This variable specifies how long a delay to allow for mouse double-clicks, in hundredths of a second. If two consecutive mouse clicks occur within the allotted time, Epsilon considers the second a double-click. Epsilon for Windows ignores this variable and uses standard Windows settings to determine double-clicks.

`draw-column-markers` Preference Default: ""

This variable may contain a series of space-separated column numbers. Epsilon for Windows draws a vertical line in the current window, at the left edge of each column number specified by this variable, counting from zero. So a value of 1 specifies a line between the first and second character positions on a line. This can be helpful when editing fixed-width files.

Set the screen-decoration color class to change the line's color.

`draw-focus-rectangle` Preference Default: 0

If nonzero, Epsilon for Windows draws a focus box around the current line, to make it easier for a user to locate the caret. A value of 1 produces a normal-sized focus rectangle.

You can customize its shape by setting this variable to a four-digit number. The four digits represent the left, right, top and bottom sides of rectangle. The digit 5 represents the normal position of that side; lower values constrict the box and higher values expand it. For instance, 5555 represents the usual box size, while 1199 represents a box that's extra narrow at its sides and extra tall.

Set the screen-decoration color class to change the box's color.

`draw-line-numbers` Preference Buffer-specific Default: 0

This variable controls whether line numbers appear to the left of each line in the current buffer. Bits in the variable control the formatting of the number. The `line-number-width` variable controls the width of the field in which line numbers will be displayed.

The 1 bit enables displaying line numbers.

The 2 bit repeats the line number when a long line is continued onto an additional screen line. Normally the line number field is blank on such continued lines.

The 4 bit positions line numbers at the left of their field. Normally they're right-aligned, ending one column before buffer text begins.

The 8 bit zero-fills the numbers, displaying 000123 instead of 123.

Epsilon omits line numbers in a pop-up or dialog window, even if the 1 bit is set, but adding the 16 bit forces line numbers there too.

`drive-match-types` Preference Default: `FSYS_LOCAL` | `FSYS_CDROM`

In Epsilon for Windows, the extended file pattern `*`: matches Windows drives like `c:` or `e:`, but only those types with a corresponding bit set in this variable. The bits defined are `FSYS_LOCAL` (16) for a local hard disk, `FSYS_NETWORK` (32) for a drive mounted from another computer, `FSYS_REMOVABLE` (64) for a flash drive or similar, and `FSYS_CDROM` (128) for a CD-ROM or similar.

`echo-line` Preference Default: 24 on a 25-line screen

This variable contains the number of the screen line on which to display the echo area, counting from zero at the top. When the screen size changes, Epsilon automatically adjusts this variable if necessary.

`eel-tab-override` Preference Default: 4

If you want the width of a tab character in EEL buffers to be different than in other buffers, set this variable to the desired value. C mode will change the buffer's tab size to the specified number of columns for EEL files (ending in `.e`).

`eel-version` Default: varies

This variable records the version number of the commands contained in the state file. Epsilon's `-quickup` flag sets this number. Epsilon compares this number to the version number stored in its executable and warns of mismatches (indicating that the state file must be updated by running `-quickup`).

`einit-file-name` Preference Default: "einit.ecm"

Epsilon loads customizations by searching for a file with this name. Commands such as `import-customizations` and `load-customizations` also use the file name in this variable to store customizations.

`epsilon-help-format-unix-gui` Preference Default: 0

This variable controls how Epsilon for Unix provides help when running as an X11 program. The value 1 makes Epsilon display help in a popup window. The values 0 and 2 make Epsilon display help in a web browser. (Epsilon for Unix always provides help in a popup window when it runs as a console program, without X11.)

`epsilon-help-format-win-console` Preference Default: 0

This variable controls how Epsilon for Windows Console provides help. The value 1 makes Epsilon display help in a popup window. The value 2 makes Epsilon display help in a web browser. The value 3 makes Epsilon display help in WinHelp. The value 0 makes Epsilon choose WinHelp on versions of Windows that always include a WinHelp program, and HTML on versions of Windows that do not (Vista and later).

`epsilon-help-format-win-gui` Preference Default: 0

This variable controls how Epsilon for Windows provides help. The value 1 makes Epsilon display help in a popup window. The value 2 makes Epsilon display help in a web browser. The value 3 makes Epsilon display help in WinHelp. The value 0 makes Epsilon choose WinHelp on versions of Windows that always include a WinHelp program, and HTML on versions of Windows that do not (Vista and later).

`epsilon-manual-port` Preference Default: 8888

When Epsilon displays its online manual in HTML format, it runs a documentation server program, and constructs a URL that tells the web browser how to talk to the documentation server. The URL includes a port number, specified by this variable. Set the variable to 0 and Epsilon won't run a local documentation server, but will instead connect to Lugaru's web site and display the manual matching Epsilon's version number. A value of -1 is similar but doesn't try to find a matching version of the manual. It uses the default copy of the manual which reflects the most recent released Epsilon version.

`errno` Default: 0

Many Epsilon primitive functions that access operating system features set this variable to the operating system's error code if an error occurs.

`expand-wildcards` Preference Default: 0

If nonzero, when you specify a file name with wild cards on Epsilon's command line, Epsilon reads each individual file that matches the pattern, as if you had listed them explicitly. If zero, Epsilon displays a list of the files that matched the pattern, in a dired buffer.

`expire-message` System Default: -1

An EEL function sometimes needs to display some text in the echo area that is only valid until the user performs some action. For instance, a command that displays the number of characters in the buffer might wish to clear that count if the user inserts or deletes some characters. After displaying text with primitives like `say()`, `note()`, or `show_text()`, an EEL function may set this variable to 1 to tell Epsilon to clear that text on the next user key.

`explicit-session-file` System Default: none

If you use the `read-session` or `write-session` commands to use a particular session file, Epsilon stores its name in this variable.

`fallback-remote-translation-type` Preference Default: 5

When you read a file that uses URL syntax, one that starts with `scp://` or `ftp://`, and Epsilon can't autodetect its line ending type, it consults this variable to determine what kind of line translation to perform. Zero indicates binary, 1 indicates DOS/Windows, 2 indicates Unix, 3 indicates Macintosh, and 5, the default, makes Epsilon for Windows assume Windows and Epsilon for Unix assume Unix.

This variable determines the translation type for new files, or files too short to permit autodetection. For most files, Epsilon will still use auto-detection. If you want to force a particular translation type for all files, bypassing auto-detection, set the `force-remote-translation-type` variable instead.

Before checking this variable, Epsilon looks for a host-specific variable. For instance, if you read a file named `ftp://example.com/filename`, Epsilon will check to see if a variable named `fallback-remote-translation-type-example-com` exists. If so, it will use that variable instead of this one. (Epsilon makes the variable name from the host name by replacing each character that isn't alphanumeric with a hyphen, then putting `fallback-remote-translation-type-` before it.)

`far-pause` Preference Default: 100

The `find-delimiter` and `show-matching-delimiter` commands pause this many hundredths of a second, when they must reposition the screen to a different part of the buffer to show the matching delimiter.

`file-date-skip-drives` Preference Default: 0

Epsilon for Windows normally warns when any file has changed on disk. Bits in this variable tell Epsilon to disable this warning for all files on certain types of devices. The value 8 disables checking on CD-ROMs. The value 4 disables checking on other removable devices like floppy disks. The value 2 disables checking files accessed over a network. The value 1 disables checking files on local hard drives. Add these values together to exclude more than one class.

`file-date-tolerance` Preference Default: 2

Epsilon warns when a file has changed on disk. Sometimes files on a network will change their recorded times shortly after Epsilon writes to them. So Epsilon ignores very small changes in a file's time. Set this variable to change the time difference in seconds that Epsilon will ignore.

`file-interpret-env-vars`

Preference

Default: 65535

Epsilon normally interprets environment variable names (and other kinds of names) in file names, using the syntax `%TEMP%` under Windows, and `$TEMP` or `${TEMP}` elsewhere. See page 125. Bits in this variable control how it operates. By default, all bits are set on.

The 1 bit permits Epsilon to substitute the value of each variable in file names you type and in `find-linked-file`. The 2 bit lets it remove any directory name appearing before such a variable name as you type a file name. The 4 bit lets it offer completion on environment variable names as you type them. You can turn these bits off if you want Epsilon to treat `%` (or `$`) as normal file name characters.

The 8 bit lets Epsilon modify the value of an environment variable which contains a list of directories, turning it into an extended file pattern that matches files in any of the directories.

The 16 bit lets Epsilon for Windows add Windows shell folders to the list of variables, each with an `@` prefix to its name to prevent collisions with environment variable names.

The 32 bit lets Epsilon load your own substitution variables from a file named `subst_vars` in your customization directory.

The 64 bit tells Epsilon to discard any environment variable or Windows shell folder whose value isn't in the form of an absolute pathname.

`file-pattern-ignore-directories`

Preference

Default: ""

In a wildcard file pattern, a directory name `**` stands for a hierarchy of directories. Epsilon recursively matches every directory within that hierarchy. This variable may be used to keep directories with certain names from being included. It contains a list of file name patterns, separated by `|` characters. Each pattern uses basic file name pattern syntax: `*` matches any number of characters, `?` matches one, and `[a-z]` specifies a character range. If a directory in the hierarchy matches this pattern, Epsilon will skip over it. For instance, `.*|*old` makes Epsilon skip over any directory ending in "old" or starting with a `.` character.

`file-pattern-rules`

Preference

Default: 0

Bits in this variable make Epsilon skip over certain types of files when interpreting file wildcard patterns. The value `FPAT_SKIP_FILE_SYMLINKS` (1) makes Epsilon ignore symbolic links to files. The value `FPAT_SKIP_DIR_SYMLINKS` (2) makes Epsilon ignore symbolic links to directories. The value `FPAT_SKIP_RECUR_SYMLINKS` (4) makes Epsilon ignore symbolic links to directories only when expanding the `**` syntax that searches entire hierarchies. Add the values to ignore multiple types of symbolic links.

Under Windows, Epsilon treats junctions and mount points the same as symbolic links to directories.

The value `FPAT_PERMIT_NO_URLS` (8) keeps file patterns that use URL syntax from matching anything.

`file-pattern-unc-domains`

Preference

Default: ""

Under Windows, file pattern matching can operate on UNC file names like `\\server\share\filename` to retrieve matching server names (when the file name pattern looks like `\\s*`, for instance). On some large local networks, it may be desirable to restrict matching to particular domains for faster completion.

If this variable is nonempty, it must contain a file name pattern, a |-separated list of domain names that can use * to match any number characters, ? to match a single character, or [a-z] to match a range of characters. Only server names from domains in this list will be retrieved. For instance, if this variable is set to EAST*|MAIN, servers in the domains EASTROOM, EASTANNEX, and MAIN would be listed, but not those in the domain WEST. (Server names aren't the same as domain names; this example setting would permit a server named \\room17, for instance, if it happened to be the EASTANNEX domain.)

Set this variable to the special setting list-domains, and Epsilon will display a list of domains on your network in a separate window the next time you complete on server names (by running dired on the pattern \\, for instance).

file-pattern-wildcards	Preference	Default: 15
------------------------	------------	-------------

Epsilon normally treats all of the characters [{};, as wildcard characters in file patterns, except when you surround a file name with "" characters. You can set this variable to force Epsilon to treat each of these characters literally. The value 1 enables comma as a wildcard character, 2 enables semicolon, 4 enables square brackets, and 8 enables curly braces. Add the values together to enable more than one group. The default of 15 enables all the above characters.

file-read-kibitz	Preference	Default: 1
------------------	------------	------------

The 1 bit in this variable tells Epsilon to display an explanatory message when it reads a file and auto-detects its translation type, but it only determines the type when it's a long way into the file. It displays a message like "filename.txt: Read as binary due to bare CR on line 303, offset 28494", which indicates the place in the file where it chose the translation type.

The 2 bit tells Epsilon to omit such messages when the grep command reads a file. The 4 bit tells Epsilon to omit such messages when file-query-replace does.

file-variables-search-limit	Preference	Default: 3000
-----------------------------	------------	---------------

When Epsilon scans the end of a file it reads, looking for file variables, the "Local Variables" line that marks the start of such variables must be within this many characters of the end of the file, or Epsilon will ignore the file's variables entirely.

filename	Buffer-specific	Default: none
----------	-----------------	---------------

This variable holds the file name associated with the current buffer.

EEL code must use assignment, not strcpy() or similar, to set it: filename = "name". It's not necessary to preserve the source character array; Epsilon will copy the value. Always set filename to an absolute pathname.

fill-c-comment-plain	Preference Buffer-specific	Default: 0
----------------------	----------------------------	------------

Set this variable nonzero if you want comment-filling commands to make C block comment lines under an initial /* start with spaces, not the usual * aligned under the initial *. (Usually this applies only to how Epsilon creates the second line of a block comment, since following lines retain the previous line's decoration.)

`fill-mode` Preference Buffer-specific Default: 0

This variable controls auto filling. If nonzero, Epsilon breaks text into lines as you type it, by changing spaces into newline characters. See the variable `c-auto-fill-mode` for the equivalent variable used in C mode buffers, and also `html-auto-fill-mode`, `xml-auto-fill-mode`, and `tex-auto-fill-mode`.

`final-macro-pause` System Default: 0

Epsilon sets this variable to 1 when a keyboard macro ends with a `pause-macro` command, to help it execute the macro correctly.

`find-lines-visible` Preference Default: 8

Epsilon uses the `find-lines-visible` variable to help determine where to position the find/replace dialog box. It considers a possible location acceptable if the top `find-lines-visible` lines of the current window can be seen behind the dialog. If fewer lines are visible, Epsilon will move the dialog to another part of the screen.

If you don't want Epsilon to ever reposition its find/replace dialog, set this variable to zero.

`find-linked-file-ignores-angles` Preference Default: 0

If this variable is nonzero, the `find-linked-file` command treats the `<>` notation in a `#include` directive found in a C/C++/Java buffer the same as the `"` notation. That is, Epsilon searches in the original file's directory for the included file, before looking in other directories. If this variable is zero, then only `#include` directives that use the `"` notation will cause Epsilon to search locally.

`first-window-refresh` Default: 1

Epsilon sets this variable prior to calling a `when_displaying` function to indicate if this is the first window the current buffer is displayed in. If 0, Epsilon has already called the `when_displaying` function for this buffer during the current screen update. Otherwise, this is the first window displaying this buffer.

`follow-mode-on` System Buffer-specific Default: 0

The `follow-mode` command uses this variable to record whether the current buffer uses follow mode. Only set it using that command.

`follow-mode-overlap` Preference Default: 0

In follow mode, set by the `follow-mode` command, a buffer displayed in multiple adjacent windows will be set so the windows always display adjacent sections of the buffer. Scrolling in one window will result in all the windows scrolling. To make it easier to see context, you may want a few lines of text at the bottom of one window repeated at the top of the next. This variable sets how many lines of overlap will be used.

`font-dialog` Preference Default: varies

This variable controls what font Epsilon for Windows uses for its dialog windows. See `font-fixed` for details on its format. Use the `set-dialog-font` command to set it.

`font-fixed` Preference Default: varies

This variable controls what font Epsilon uses.

Under Windows, it contains the name of the font, followed by several numbers separated by commas. The first number specifies a point size. The second specifies the width of the font in pixels, or 0 if any width is acceptable. A small number of fonts, such as Terminal, have multiple widths for each height. The third number specifies how bold the font is. A typical font uses a value of 400, while a bold font is usually 700. The fourth number is nonzero to indicate an italic font. The fifth number indicates a character set; 1 means use the default character set for the font, 0 means use ANSI, 255 means use OEM.

Under X11, it contains a standard X11 font name, possibly with wildcards.

EEL code must use assignment, not `strcpy()` or similar, to set this and other font variables, as in `font_fixed = "Courier New,12,0,400,0,0"`. It's not necessary to preserve the source character array; Epsilon will copy the value. Epsilon automatically saves and restores font variables using OS-specific methods (in the Windows registry, or in an X11 resource file), and does not save these in its state file.

`font-printer` Preference Default: varies

This variable controls what font Epsilon for Windows uses when printing. See `font-fixed` for details on its format. Use the `set-printer-font` command to set it.

`font-styles` Preference Default: 7

Epsilon normally uses various styles of the selected font (italic, bold, underlined) for particular text, such as comments. Bits in this variable say which of these styles it can use. The value 1 permits a bold style, 2 permits italic, and 4 permits underlined text. Set this variable to the sum of the permitted styles; a value of 0 disables all font styles.

`font-styles-tolerance` Preference Default: 1

When Epsilon for Windows looks for matching styles of the selected font (italic, bold, or underlined), it ignores any that are too different in size from the base font. This variable determines how different such fonts may be. It must be a number from 0 to 999. The rightmost digit (ones digit) sets the permitted variation in height, and the middle digit (tens) sets the permitted variation in width: 0 requires an exact match. The leftmost digit (hundreds) influences the permitted additional overhang/underhang of characters, rejecting fonts that are much more slanted than the base font.

If Epsilon shows no bold, italic, or underlined text with your selected font, you can try setting these values higher; higher values permit less attractive font combinations, though.

Epsilon for Unix requires font styles to match the original font exactly, always acting as if this variable were zero.

`force-remote-translation-type` Preference Default: 5

When you read a file that uses URL syntax, one that starts with `scp://` or `ftp://`, Epsilon consults this variable to determine what kind of line translation to perform. Zero forces binary, 1 forces DOS/Windows, 2 forces Unix, 3 forces Macintosh, and 5, the default, lets Epsilon autodetect the file type.

auto-detection isn't supported. It may also be useful when your connection to the host is very slow. With this option, whenever you type a directory name, you must either end it with a / character to indicate it's a directory, or use the `dired` command instead of one of the other file-reading commands.

`ftp-passive-transfers` Preference Default: 1

This variable controls how Epsilon's FTP client transfers files. Epsilon knows three methods, called "passive", "active", and "default port". Firewalls or ancient FTP server software can cause one or more of the methods to fail. Set this variable to zero to use only active transfers. Set it to two to make Epsilon try active transfers first, then passive. Set it to three to make Epsilon use the "default port" method. The default of one makes Epsilon try passive, then active.

`full-key` Default: none

This variable holds the value of the last key pressed, or a special code indicating a mouse event. For keys that have both a generic and specific interpretation, such as `<Backspace>` (which is generically `Ctrl-h`), this variable holds the specific interpretation, while `key` has the generic one. For other keys, these two variables are equal.

`full-path-on-mode-line` Preference Default: 0

Set this variable to 1 if you want Epsilon to display the full path of each file on the mode line. The default value of 0 makes it use a path relative to the current directory (set by the `cd` command) whenever it can. The value 2 makes it always display only the last component of the file name (its base name).

`full-redraw` Default: 0

If nonzero, Epsilon rebuilds all mode lines, as well as any precomputed information Epsilon may have on window borders, screen colors, and so forth, on the next redisplay. Epsilon then resets the variable to zero.

`fundamental-auto-show-delim-chars` Default: ""

This variable holds the set of delimiter characters that should trigger Epsilon's auto-show-delimiters feature in Fundamental mode. Epsilon will search for and highlight the match of each delimiter.

Delimiters in the left half of the list must be left-delimiters and those in the right half must be right-delimiters, as in `([])`.

`fundamental-spell-options` Preference Default: 0

This variable controls the Spell minor mode in Fundamental mode. Use the `spell-mode` command to set it to 1, and Epsilon will highlight misspelled words in Fundamental mode. See the `default-spell-options` variable for the other bits you can set to customize spell checking in Fundamental mode.

`fwd-search-key` Preference Default: -1

Inside a search command, Epsilon recognizes a key with this key code as a synonym for Ctrl-s, for pulling in a default search string or changing the search direction.

`gams-auto-show-delim-chars` Default: "[()]"

This variable holds the set of delimiter characters that should trigger Epsilon's auto-show-delimiters feature in GAMS mode. Epsilon will search for and highlight the match of each delimiter.

`gams-files` Preference Default: 0

The file extensions .inc, .map, and .dat are used in the GAMS language for mathematical programming. But they're also commonly used to represent other things. By default Epsilon assumes such files are not GAMS files; set this variable nonzero if you want Epsilon to assume they are GAMS files.

`global-spell-options` Preference Default: 0xf

Bits in this variable control certain spell checking options:

0x01 (1): This bit makes Epsilon automatically save an ignore list file after adding an entry.

0x02 (2): This bit makes Epsilon delete a file-specific ignore list file when you delete its main file.

0x10 (16): This bit makes Epsilon display an option for adding a misspelled word to an extension-specific ignore list. You can enable this if you prefer to base ignore lists on file extensions, not just modes. Epsilon looks for extension-specific ignore list files regardless of this bit, but doesn't offer to add to them without it.

`goal-column` Buffer-specific Default: -1

If the `goal-column` variable is non-negative, the up-line and down-line commands and commands that move by pages always move to the goal column. If `goal-column` is negative, the commands try to remain in the same column. When a region is highlighted, Epsilon ignores the goal column.

`got-bad-number` System Default: 0

Several EEL functions that convert a character string into a number set this variable to indicate whether the string held a valid number.

`grep-default-directory` Preference Default: 0

When you press <Enter> without entering a file pattern for `grep`, and this variable is zero (the default), Epsilon tries to search the same set of files as last time, even if you've subsequently changed directories.

Set this variable to 1 if you want `grep` to instead reinterpret the file pattern you typed according to the current directory.

Set it to 2 if you want Epsilon to reinterpret the previous file pattern according to the directory associated with the current buffer. Set it to 3 if you want Epsilon to interpret any relative pattern you type according to the directory associated with the current buffer.

For settings 2 and 3 above to work properly, the `grep-prompt-with-buffer-directory` variable must not be set to 0.

This variable affects only relative file patterns like `*.cpp`. If you typed an absolute path the first time, pressing (Enter) will always search those files regardless of this setting.

`grep-empties-buffer` Preference Default: 0

By default, each invocation of `grep` appends its results to the `grep` buffer. If you set the variable `grep-empties-buffer` to a nonzero value, `grep` will clear the `grep` buffer at the start of each invocation.

`grep-ignore-file-basename` Preference Default: ""

This variable contains a regular expression pattern. The `grep` and `file-query-replace` commands skip over file names whose base names (the part after the last `/` or `\`) start with text that matches this pattern. If the variable is zero-length, the commands ignore it.

`grep-ignore-file-extensions` Preference Default:
"|.obj|.exe|.o|.b|.dll|... (omitted) |"

This variable contains a list of file name extensions for `Epsilon` to skip over during a `grep` or `file-query-replace` command. Each extension must appear surrounded by `'|'` characters.

`grep-ignore-file-max-size` Preference Default: 2,000,000,000

The `grep` and `file-query-replace` commands skip over files bigger than this size (on disk). A value of 0 or negative removes the size limit, so the commands will search files of any size.

`grep-ignore-file-pattern` Preference Default:
"((.*[\\/])?(CVS|.svn) [\\/] |/proc/"

This variable contains a regular expression pattern. The `grep` and `file-query-replace` commands skip over file names whose absolute forms start with text that matches this pattern. If the variable is zero-length, the commands ignore it.

`grep-ignore-file-types` Preference Default: 7

Bits in this variable tell the `grep` and `file-query-replace` commands to skip over files of certain types. The value 1 makes the commands ignore file names that refer to devices. The value 2 ignores pipes, and the value 4 ignores any files that are neither ordinary files nor one of the above types.

`grep-include-timestamp` Preference Default: 0

If nonzero, the `grep` command includes the file's timestamp at the start of each result line in the `grep` buffer. A value of 1 uses month names; a value of 2 results in an all-numeric representation suitable for sorting.

`grep-keeps-files` Preference Default: 0

If nonzero, the `grep` command reads each file matching the supplied pattern using the `find-file` command. If zero, `Epsilon` reads each file into a temporary buffer and discards the buffer after it finishes listing the matches.

grep-on-changed-file

Preference

Default: 1

This variable sets what happens when the `grep` command notices a buffer's file has been changed on disk.

If 0, Epsilon prompts, asking if it should reread the file, replacing the version in memory, or ignore the file and only search the version in memory.

If 1, Epsilon never prompts, and ignores the version on disk, adding a warning in the `grep` buffer that only the version in memory was searched. (The `locate-file` on `Ctrl-x Ctrl-l` is a convenient way to get from the `grep` buffer to the file's buffer.)

If 2, Epsilon automatically reads the modified file from disk, as long as the version in memory hasn't been edited. If it's been edited, Epsilon prompts.

If 3, Epsilon rereads the file if the buffer hasn't been modified; otherwise it searches the version in memory and adds a warning to the `grep` buffer. It doesn't prompt in either case.

grep-prompt-with-buffer-directory

Preference

Default: 1

The `grep-prompt-with-buffer-directory` variable controls how the `grep` and `file-query-replace` commands use the current directory at file prompts.

When this variable is 1, the default, Epsilon temporarily changes to the current buffer's directory while prompting for a file name, and interprets file names relative to that directory. It also arranges for each mode line to display the full path of its file while prompting, to make it easier to see which directory will be used.

When this variable is 2, Epsilon inserts the current buffer's directory at the prompt. You can type a relative file pattern like `*.cpp`. Or you can begin typing a new absolute file pattern like `/etc/*.cf` or `c:\windows*.ini` right after the inserted pathname. Epsilon will delete the inserted pathname when it notices your absolute pathname.

A setting of 3 makes Epsilon insert the current buffer's directory in the same way, but keeps Epsilon from automatically deleting the inserted pathname if you type an absolute one.

When this variable is 0, relative file names will be interpreted based on the global current directory set with the `cd` command, not any directory associated with the current buffer.

If you press `(Enter)` without typing anything, and a directory used in the above descriptions has changed, the `grep-default-directory` variable controls whether Epsilon reinterprets the previous file pattern according to the new directory, or searches the same set of files as before. See that variable for more details.

grep-show-absolute-path

Preference

Default: 0

This variable controls how the `grep` command formats the file names it inserts into the `grep` buffer for each match. If `grep-show-absolute-path` is 0, Epsilon uses a relative pathname whenever it can. If 1, Epsilon uses an absolute pathname always. If 2, Epsilon for Windows lists each file with its absolute DOS-style 8+3 file name. (This setting is the same as 1 in environments without such a notion.)

gui-cursor-shape

System

Default: 100002

This variable holds the current cursor shape code under Windows and Unix's X11 windowing system. Epsilon copies values from `overwrite-gui-cursor`, `normal-gui-cursor`, or one of the other cursor variables, as appropriate, into this variable whenever you switch windows or

buffers. Set those variables instead of this one. Epsilon only uses this variable under Windows and X. See `cursor-shape` for the non-graphical equivalent.

The X11 version of Epsilon can only change the cursor shape if you've provided an `Epsilon.cursorstyle:1` resource (see page 6).

`gui-menu-file` Preference Default: "gui.mnu"

This variable contains the name of the file Epsilon loads its menu from at startup, in the Windows version.

`has-arg` Default: 0

Epsilon indicates that a command has received a numeric argument by setting this variable nonzero. The value of the numeric argument is in the `iter` variable.

`has-feature` Default: varies

Epsilon runs under various operating systems. Some OS versions of Epsilon have a few features that others lack. An EEL function may test bits in this variable to check if certain features are available.

`hex-overtypemode` Preference Default: 0

Set this variable to one if you want hex mode to begin in its overtype submenu. See `hex-mode`.

`html-asp-coloring` Preference Default: 1

This variable tells Epsilon how to syntax highlight scripts embedded in `<% %>` delimiters in HTML documents, when the file doesn't name any specific script language. Zero means use a single color, 1 means color as Javascript, 2 means color as VBScript, 3 means color as PHP, 4 means Python, 5 means CSS, and 10 means `<%` doesn't start an embedded script.

`html-auto-fill-combine` Preference Default: 10

When auto-filling breaks a line of HTML, it avoids breaking a line that contains just the start of an element, except when the element name is longer than specified by this variable. For example, it won't split a line right after `<a` at its start, but it may after `<customquote`.

`html-auto-fill-mode` Preference Default: 31

This variable controls whether Epsilon automatically breaks long lines as you type in HTML mode and the related PHP mode. The 1 bit toggles filling on and off entirely. It's set by the `auto-fill-mode` command. Other bits control where filling occurs.

The 2 bit lets it break text that isn't part of an HTML tag. The 4 bit lets it break HTML tags. The 8 bit lets it break HTML comments. The value 16 lets it break comments in any embedded scripting. By default, Epsilon breaks in all these regions.

`html-auto-indent`

Preference

Default: 0xff

This variable controls automatic indentation when you press `<Enter>` in HTML mode. Bits in the variable control whether Epsilon auto-indents in specific regions of the document. The 0x1 bit makes Epsilon auto-indent when you press `<Enter>` outside script blocks. Other bits, as shown in the table below, make Epsilon auto-indent in that type of scripting.

Scripting Language	Bit
JavaScript	0x2
VBScript	0x4
PHP	0x8
Python	0x10
CSS	0x20

When you disable smart auto-indenting in a certain type of scripting by setting one of these bits, `<Enter>` will instead indent to the previous line. Set the 0x4000 bit if you prefer no indenting at all for regions where you've disabled smart indenting.

`html-auto-show-delim-chars`

Default: "<>"

This variable holds the set of delimiter characters that should trigger Epsilon's auto-show-delimiters feature in HTML mode. Epsilon will search for and highlight the match of each delimiter.

`html-display-definition`

Preference

Default: 15

Bits in this variable control whether Epsilon displays the name of the current function within HTML, XML, or PHP scripting, or HTML or XML nesting elements outside of scripting. The value 1 enables displaying the current function name in scripting. The value 2 makes Epsilon display the type of scripting, such as "JavaScript" or "VBScript", on scripting lines that are outside of any function. The 4 bit makes Epsilon show the innermost elements in effect at the start of the current line in XML mode. The 8 bit is similar, but applies to HTML mode. The `html-display-nesting-width` variable influences how many elements appear.

`html-display-nesting-width`

Preference

Default: 40

In HTML and XML modes, the mode line displays the names of the innermost tags in effect at the start of the current line. This variable controls how much space is allocated for that display. Smaller values let Epsilon use more of the mode line for tag names, so the buffer's file name is more likely to be abbreviated. You can use the `html-list-element-nesting` command to display all the tags.

`html-empty-elements`

Default:

`"|hr|isindex|col|... (omitted) |"`

HTML mode treats elements named in this list as empty (having no end tag). Each element name must be surrounded by `'|'` characters.

`html-indent`

Preference Buffer-specific

Default: 3

Each level of indentation in HTML mode will occupy this many columns. If this variable is zero, Epsilon uses the tab size instead.

`html-indenting-rules`

Preference

Default: `0xffffffff`

This variable controls whether Epsilon uses smart indenting for HTML or XML (or scripting embedded within it), or just copies indentation from the previous line. Bits in the variable enable smart indenting for various languages, as shown below.

Scripting Language	Smart Indent
JavaScript	0x1
VBScript	0x2
PHP	0x4
Python	0x8
CSS	0x10

These bits enable indenting in certain contexts:

0x100 permits indenting inside tags.

0x200 permits indenting for text lines that don't start with a tag.

0x400 permit indenting the first line of a script block, but only when its opening `<` is indented. When the `<` is at the left margin, the script will start there too.

0x800 permit indenting the first line of a script block unconditionally, overriding 0x400.

Also see the `html-auto-indent` and `xml-auto-indent` variables, which control whether and how Epsilon indents when you press `(Enter)` (not `(Tab)` or other keys that indent).

By default smart indenting is enabled in all cases.

When there are multiple start tags on a single line, Epsilon normally indents the next line by one level no matter how many start tags appear. Turning off the 0x80 bit tells Epsilon to indent more in such cases, as if the start tags had appeared on separate lines.

The 0x8000 bit controls how Epsilon indents when the attributes of a tag continue onto the next line. This bit makes Epsilon align the equal signs following each attribute. Turn it off if you want Epsilon to align the attribute names instead.

When you press certain keys like `{`, `:`, `#`, or a right delimiter, Epsilon normally reindents that line when in scripting. These bits in `html-indenting-rules` enable such reindenting:

Scripting Language	Enable Reindent
JavaScript	0x10000
VBScript	0x20000
PHP	0x40000
Python	0x80000
CSS	0x100000

Note that you can type expressions at the `set-variable` prompt, like `varname | 0x100` to set a bit in the variable `varname` you're setting, or `varname & ~ 0x100` to turn it off.

`html-no-indent-elements`

Preference

Default: `"|html|body|"`

This variable contains a list of HTML element names, surrounded by and separated by `|` characters. These elements won't receive additional indenting. Epsilon also assumes the end tags of these elements are optional.

html-other-coloring Preference Default: 1

This variable tells Epsilon how to syntax highlight scripts marked with an unknown language name embedded in HTML documents. Zero means use a single color, 1 means color as Javascript, 2 means color as VBScript, 3 means color as PHP, 4 means Python, 5 means CSS, and 10 means such tags don't start an embedded script.

html-paragraph-elements Preference Default:
"p|html|head|... (omitted)"

When filling paragraphs in HTML mode, Epsilon treats tags matching element names on this list as paragraph delimiters. The list uses the syntax of a regular expression pattern.

html-paragraph-is-container Preference Buffer-specific Default: 1

In modern HTML, the p element for paragraphs is used as a container, with <p> at the start of a paragraph and </p> at its end. In very old HTML, the p tag is used as a paragraph separator, appearing between paragraphs, and </p> isn't used. For correct indenting and similar purposes, Epsilon needs to know which method the current file uses.

A value of 1, the default, tells Epsilon the file uses the modern method where <p> begins a new indenting level that </p> will close later. A value of 0 tells Epsilon the file uses the older method where <p> is a mere separator and causes no additional indentation. A value of 2 tells Epsilon to examine the file, the first time it needs to decide this, and guess at which method it uses.

html-php-coloring Preference Default: 3

This variable tells Epsilon how to syntax highlight scripts embedded in <? ?> delimiters in HTML documents, when the file doesn't name any specific script language. Zero means use a single color, 1 means color as Javascript, 2 means color as VBScript, 3 means color as PHP, 4 means Python, 5 means CSS, and 10 means <% doesn't start an embedded script.

html-prevent-coloring Preference Default: 0

Bits in this variable prevent language-specific syntax highlighting of embedded scripting in HTML or XML mode, by forcing Epsilon to ignore the sequence that begins that type of scripting.

Each bit in the table below makes Epsilon ignore the corresponding scripting type (and certain equivalent constructions as well).

Bit	Scripting Tag
0x1	<script language=jscript>
0x1	or <script language=javascript>
0x1	or <script language=ecmascript>
0x2	<script language=vbscript>
0x4	<script language=php>
0x4	or <?php>
0x8	<script language=python>
0x8	or <?python>
0x10	<jsp:scriptlet>
0x20	<cfscript>

`html-recognize-coldfusion-comments` Preference Buffer-specific Default: 1

If nonzero, HTML mode interpret the sequence `<!--` as the start of a ColdFusion comment, terminated by a `-->` and allowing nested comments. If zero, it follows normal HTML comment syntax rules, with no nesting.

`html-reindent-previous-line` Preference Default: 1

This variable controls whether Epsilon reindents the previous line when you press `<Enter>` in HTML mode.

`html-spell-options` Preference Default: 0

This variable controls the Spell minor mode in HTML mode. Use the `spell-mode` command to set it to 1, and Epsilon will highlight misspelled words in HTML text, ignoring element names and attributes. See the `default-spell-options` variable for the other bits you can set to customize spell checking in HTML mode.

`html-style-rules` Preference Buffer-specific Default: 0

This variable contains bits that modify how Epsilon interprets HTML tags.

The 1 bit tells Epsilon that each start tag that has no matching end tag will be marked as such using the `/>` syntax for self-terminating tags, as in XHTML. Epsilon will assume for indenting and other purposes that a tag like `<hr>` begins a block that a later `</hr>` will end; the self-terminating form `<hr />` doesn't begin such a block.

Without this bit, HTML mode uses the `html-empty-elements`, `coldfusion-empty-elements`, and `html-no-indent-elements` variables, as well as built-in logic for the elements `li`, `option`, `p`, `dd`, and `dt`, to decide which elements should have no end tags.

`html-tag-match-look-back` Preference Default: 10,000,000

With some HTML start tags like `<dt>`, end tags are optional. Sometimes Epsilon must scan the buffer looking for an end tag to determine correct indenting. It restricts its search to the number of characters specified by this variable.

`http-log-request` Preference Default: 0

If nonzero, accessing an http URL includes the actual http request Epsilon sent in the "HTTP Headers" buffer, just before the response.

`http-proxy-exceptions` Preference Default: "|localhost|127.0.0.1|"

When Epsilon uses a proxy server, it still directly connects to host names in this list. Each entry must have a `—` character before and after.

`http-proxy-port` Preference Default: 0

If you want Epsilon to use a proxy server to retrieve web pages, set its port number here, and set `http-proxy-server` to the proxy server's name. Zero means no proxy.

<code>http-proxy-server</code>	Preference	Default: ""
--------------------------------	------------	-------------

If you want Epsilon to use a proxy server to retrieve web pages, set its name here, and set `http-proxy-port` to the appropriate port setting.

<code>http-user-agent</code>	Preference	Default: ""
------------------------------	------------	-------------

When Epsilon retrieves a web page in response to an http URL, it identifies itself to the web server as “Epsilon *versionnumber*”. Set this variable to force Epsilon to use a different name.

<code>idle-coloring-delay</code>	Preference Buffer-specific	Default: 100
----------------------------------	-------------------------------	--------------

When Epsilon isn’t busy acting on your keystrokes, it looks through the current buffer and assigns colors to the individual regions of text, so that Epsilon responds faster as you scroll through the buffer. For smoother performance, Epsilon doesn’t begin to do this until it’s been idle for a certain period of time, specified by this variable. Set it to the number of hundredths of a second to wait before computing more coloring information. With its default value of 100, Epsilon waits one second. Set it to -1 to disable background code coloring.

<code>idle-coloring-size</code>	System Buffer-specific	Default: 1000
---------------------------------	---------------------------	---------------

While waiting for the next keystroke, Epsilon syntax-highlights the rest of the current buffer to improve performance. It highlights in small sections. This variable determines the size of each section. Some language modes highlight faster when they can work with larger sections.

<code>ignore-error</code>	Preference	Default: none
---------------------------	------------	---------------

This variable holds a regular expression that commands like `next-error` use to filter out any error messages Epsilon should skip over, even if they match the error pattern. For example, if `ignore-error` contains the pattern “`.*warning`”, Epsilon will skip over any error messages that contain the word “warning”.

<code>ignore-file-basename</code>	Preference	Default: "%."
-----------------------------------	------------	---------------

This variable contains a regular expression pattern. File name completion skips over file names whose base names (the part after the last / or \) start with text that matches this pattern. If the variable is zero-length, completion ignores it.

<code>ignore-file-extensions</code>	Preference	Default: ".obj .exe .o .b "
-------------------------------------	------------	-----------------------------

This variable contains a list of file name extensions for Epsilon to ignore during file name completion. Each extension must appear surrounded by ‘|’ characters. If no matches result when using this variable, completion ignores it and tries again.

<code>ignore-file-pattern</code>	Preference	Default: ""
----------------------------------	------------	-------------

This variable contains a regular expression pattern. File name completion skips over file names whose absolute forms start with text that matches this pattern. If the variable is zero-length, file name completion ignores it.

`ignore-kbd-macro` Default: 0

When the `ignore-kbd-macro` variable is nonzero, Epsilon suspends any running keyboard macros and doesn't retrieve keys from them. When zero (the default), Epsilon retrieves keys from a keyboard macro before handling keys from the keyboard.

`ignoring-file-change` System Buffer-specific Default: 0

Epsilon sets this variable nonzero when the user says to temporarily ignore file date warnings. See `want-warn`.

`in-echo-area` Default: 0

The `in-echo-area` variable controls whether the cursor is positioned at point in the buffer, or in the echo area at the bottom of the screen. The `sayput()` primitive sets this variable, `say()` resets it, and it is reset after each command.

`in-perl-buffer` System Buffer-specific Default: 0

Epsilon's C mode uses this to record if the current buffer is really in Perl mode (which is implemented as a variant of C mode).

`in-shell-buffer` System Buffer-specific Default: 0

Epsilon's Perl mode uses this to record if the current buffer is really in Shell mode (which is implemented as a variant of Perl mode).

`include-directories` Preference Default: ""

The `find-linked-file` command, in C/C++/Java buffers, edits the file named by the `#include` directive on the current line. Epsilon knows a few standard places to look for `#include` files, but if Epsilon doesn't find yours, set this variable to a list of directories where Epsilon should look, in addition to the standard places. Separate the directory names with colons under Unix, with semicolons elsewhere.

`indent-comment-as-code` Preference Default: 1

If nonzero, commenting commands indent lines containing only a comment to the same indentation as other text.

`indent-preprocessor-contin` Preference Buffer-specific Default: 4

Epsilon indents continued preprocessor lines based on the indentation established by the initial line being continued. If that line doesn't have enough code to establish an indentation, Epsilon uses the indentation specified by this variable. If the variable is zero, Epsilon uses simplified indenting rules for preprocessor continuation lines.

`indent-with-tabs` Preference Buffer-specific Default: 1

If zero, Epsilon indents using only space characters, not tab characters.

`indents-separate-paragraphs` Preference Buffer-specific Default: 0

Blank lines and `^L` characters always separate paragraphs. If the variable `indents-separate-paragraphs` has a nonzero value, then a paragraph also begins at a nonblank line that starts with a tab or a space.

`info-path-non-unix` Preference Default: `%x..\info;%x`

Epsilon's info mode looks for Info files in each of the directories named by this variable (but see `info-path-unix` for the Unix equivalent). Separate the directory names with semicolons. The sequence `%x` tells Epsilon to substitute the directory containing its executable.

`info-path-unix` Preference Default: (omitted)

Under Unix, Epsilon's info mode looks for Info files in each of the directories named by this variable. (See `info-path-non-unix` for the non-Unix equivalent). Separate the directory names with colons. The sequence `%x` tells Epsilon to substitute the directory containing its executable.

`info-recovering` System Default: 0

Epsilon's info mode uses this variable internally to record whether it's currently recovering from failing to reach a missing node.

`initial-tag-file` Preference Default: `"default.tag"`

This variable holds the name of the tag file Epsilon will search for. If it holds a relative pathname, Epsilon will search for the file in the current directory tree. If `initial-tag-file` holds an absolute pathname, Epsilon will always use that tag file.

`insert-default-response` Preference Default: 1

If this variable is 1, at many prompts Epsilon will insert a default response before you start typing. The default response will be highlighted, so typing any text will remove it. If you turn off `typing-deletes-highlight`, you may wish to set this variable to 0.

While prompting for text, Epsilon can temporarily set this variable to other values. A value of 2 makes Epsilon insert its default text without highlighting it. This means the text won't automatically be deleted when you begin typing. A value of 3 inserts the default text, doesn't highlight it, and prepares to modify your file name response as you type it. See the description of `prompt-with-buffer-directory`.

`insert-file-remembers-file` Preference Default: 0

Set this nonzero and the `insert-file` and `write-region` commands will prompt with the name of the last inserted or written file as the default. Set it to zero and they'll offer the current buffer's directory as a default.

`invisible-window` System Window-specific Default: 0

If nonzero, Epsilon won't display the text of the window (although it will display the border, if the window has one). Epsilon won't modify the part of the screen that would ordinarily display the window's text.

<code>is-current-window</code>	System	Default: none
<p>An EEL program may set a highlighted region to be controlled by this variable to signal that the region should only be displayed in the current window, not in other windows that display the same buffer.</p>		
<code>is-gui</code>		Default: varies
<p>The <code>is-gui</code> variable indicates whether a version of Epsilon with native dialogs is running. In this version it's set to 2 (<code>IS_NT</code>) in Epsilon for Windows, to 0 in Epsilon for Windows Console and all non-Windows versions.</p>		
<code>is-mswindows</code>		Default: varies
<p>This variable is nonzero if a version of Epsilon for Windows is running, either the GUI version or the Windows Console version. The constant <code>IS_MSWINDOWS_64BIT</code> means the 64-bit version of Epsilon is running. The constant <code>IS_MSWINDOWS_32BIT</code> means the 32-bit version of Epsilon is running.</p>		
<code>is-unix</code>		Default: varies
<p>This variable is nonzero if Epsilon for Unix is running. It's set to the constant <code>IS_UNIX_XWIN</code> if Epsilon is running as an X11 program, or <code>IS_UNIX_TERM</code> if Epsilon is running as a curses program.</p>		
<code>is-unix-flavor</code>		Default: varies
<p>This variable is nonzero if Epsilon for Unix is running. It's set to the constant macro <code>IS_UNIX_LINUX</code> in the Linux version, <code>IS_UNIX_BSD</code> in the FreeBSD version, and <code>IS_UNIX_MACOS</code> in the Mac OS version.</p>		
<code>iter</code>		Default: 1
<p>Epsilon indicates that a command has received a numeric argument by setting the <code>has-arg</code> variable nonzero, and setting <code>iter</code> to the value of the numeric argument.</p>		
<code>java-indent</code>	Preference Buffer-specific	Default: 4
<p>C mode indents Java files by this many columns for each additional level of nesting. If the variable is less than or equal to zero, Epsilon uses the value of <code>tab-size</code> instead. Set this variable nonzero if you want Epsilon to use one number for displaying tab characters, and a different number for indenting Java code. (Epsilon will indent using a combination of spaces and tabs, as necessary.)</p>		
<code>jump-to-dvi-command</code>		Default: ""
<p>This variable contains the command line that the <code>jump-to-dvi</code> command should use. See page 98 for more information. Epsilon expands these sequences in this variable to form the command line:</p>		
<p><code>%l</code> Epsilon substitutes the TeX source file's line number.</p>		

%f Epsilon substitutes the full path of the TeX source file.

%d Epsilon substitutes the full path of the DVI or PDF file.

%b Epsilon substitutes the base name of the TeX source file, not including any path or extension.

%p Epsilon substitutes the name of the directory containing the TeX source file, not including any path or extension.

`jump-to-dvi-extension` Default: ".dvi"

The `jump-to-dvi` command can operate with either .dvi or .pdf files. Set this variable to either “.dvi” or “.pdf” depending on which kind of viewer you want to use to view your TeX output.

`jump-to-dvi-main-file` Default: ""

This variable contains the name of the .dvi or .pdf file that `jump-to-dvi` command should use. Running `jump-to-dvi` with a prefix argument prompts to set it. See that command for details. It's empty if the command should look for a .dvi or .pdf file based on the name or contents of the current TeX file each time it's run.

`key` Default: none

This holds the value of the last key pressed, or a special code indicating a mouse event. For keys that have both a generic and specific interpretation, such as `<Backspace>` (which is generically `Ctrl-h`), this variable holds the generic interpretation, while `full-key` has the specific one. For other keys, these two variables are equal.

`key-code` Default: none

This variable contains the sixteen-bit BIOS encoding for the last key that Epsilon received from the operating system. Its ASCII code is in the low eight bits and its scan code is in the high eight bits. This variable is always 0 when Epsilon can't get this information, or under Unix.

`key-from-macro` Default: varies

This variable is nonzero whenever the most recent key (or mouse event) came from a keyboard macro, not an actual keypress.

`key-is-button` System Default: varies

When you click on a button in a dialog, Epsilon returns a particular fixed keystroke: `Ctrl-m`, the abort key specified by the `abort-key` variable, or the help key specified by the `HELPKEY` macro. To distinguish actual keys from these buttons, Epsilon sets the `key_is_button` variable to zero when returning normal keys, and to a nonzero value when returning one of these button keys. It uses a specific value for each button in a dialog.

`key-repeat-rate` Preference Default: 40

Under Windows, this variable controls the rate at which keys repeat in Epsilon, in repeats per second. Setting this variable to 0 lets the keyboard determine the repeat rate, as it does outside of Epsilon. Setting this variable to -1 makes keys repeat as fast as possible. Epsilon never lets repeated keys pile up; it automatically ignores repeated keys when necessary.

`key-type` Default: none

This variable has a special code that identifies the type of key pressed. Epsilon uses the key type to implement its auto-quoting facility.

`kill-buffers` Preference Default: 10

This variable holds the maximum number of kill buffers, for holding killed text. Setting this variable to a new value makes Epsilon throw away the contents of all the kill buffers the next time you execute a command that uses kill buffers.

`kill-rectangle-removes` Preference Default: 0

This variable controls the default behavior of the `kill-rectangle` command. If zero, it replaces the rectangular block with spaces. If nonzero, it removes the columns of the rectangular block entirely, shifting text leftward, like the `delete-rectangle` command. No matter which behavior is set by this variable, running `kill-rectangle` with a numeric prefix argument produces the other behavior.

`last-index` Default: none

The `do_command()` primitive copies the name table index it receives as a parameter into this variable, just before it executes the indicated command, so the help system can find the name of the current command (among other uses).

`last-show-spaces` System Buffer-specific Default: 0

Epsilon records the previous value of the `show-spaces` variable here, to detect changes to it.

`last-window-color-scheme` System Default: 0

When Epsilon has been set to display its tiled windows without borders (via the `toggle-borders` command), it uses this variable to help assign separate color schemes to the individual windows.

`latex-2e-or-3` Preference Default: 1

Set this variable to 0 if you want LaTeX mode commands like `tex-italic` on Alt-i to insert a LaTeX 2.09 command, instead of a LaTeX 2e/3 command. (For example, `tex-italic` inserts `\textit` in LaTeX 2e/3 mode, and `\it` otherwise.)

`latex-display-math-env-pat` Preference Default:
"equation|displaymath|eqnarray%*?"

This variable holds a regular expression pattern that matches environment names whose contents Epsilon should color using the `tex-display-math` color class.

`latex-math-env-pat` Preference Default: "math"

This variable holds a regular expression pattern that matches environment names whose contents Epsilon should color using the `tex-math` color class.

<code>latex-non-text-argument</code>	Preference	Default: " <code> begin end label ... (omitted) </code> "
--------------------------------------	------------	---

Epsilon uses the `tex-keyword` color class (instead of `tex-text`) for the argument of these LaTeX commands. This keeps the spell checker from checking these environment names, labels, and cite tags. Surround each command name with `|` characters. You can modify the list if you want spell checking to check such tags.

<code>latex-spell-ignore-pattern-prefix</code>	Preference	Default: (omitted)
--	------------	--------------------

Spell mode ignores a word in LaTeX mode when the text just before it matches this pattern. The default pattern makes it ignore environment names, labels, and cite tags, even in contexts like comments where they aren't colored as keywords.

<code>latex-spell-options</code>	Preference	Default: 0
----------------------------------	------------	------------

This variable controls the Spell minor mode in LaTeX mode. Use the `spell-mode` command to set it to 1, and Epsilon will highlight misspelled words in LaTeX text. See the `default-spell-options` variable for the other bits you can set to customize spell checking in LaTeX mode.

Also see the `latex-non-text-argument` variable.

<code>latex-tag-keywords</code>	Preference	Default: " <code>label</code> "
---------------------------------	------------	---------------------------------

This variable contains a list of LaTeX keywords, separated by `|` characters. Tagging commands like `tag-files` and `list-definitions` treat the argument of any command in this list as a tag.

<code>leave-blank</code>		Default: 0
--------------------------	--	------------

When Epsilon is about to exit, it normally redisplay each mode line one last time just before exiting, but only if this variable is zero.

<code>line-in-window</code>		Default: none
-----------------------------	--	---------------

On each screen refresh, Epsilon sets this variable to the line of point within the current window, counting from zero. If you switch windows or move point, Epsilon will not update this variable until the next refresh.

<code>line-number-width</code>	Preference Buffer-specific	Default: 6
--------------------------------	----------------------------	------------

Set the `draw-line-numbers` variable to have Epsilon display a line number to the left of each line. It fits the line number into the number of columns specified by this variable. When a number won't fit, Epsilon displays a `*` in the first column and shows part of the number.

<code>list-definitions-live-update</code>	Preference	Default: 1
---	------------	------------

Bits in this variable control how the `list-definitions` dialog displays its results. The value 1 makes it reposition in the original window as you move from one definition to the next, showing each definition line in context. The value 2 makes it include the source file line number of each definition it displays in the window's title.

`list-which-definitions` Preference Default: 0xff

Bits in this variable determine which items are included in the list of definitions produced by the `list-definitions` command. The value 0x1 includes functions, 0x2 includes variables, 0x4 includes macros, 0x8 includes the values of enums, 0x10 includes the names of classes, structures, unions, and similar items, and 0x20 includes the names of enums. Also see `tag-which-items`. Some modes may not distinguish all these types of items.

`load-customizations` Preference Default: 1

This variable sets whether Epsilon reads its `einit.ecm` customization file when it starts. A value of 1 makes Epsilon search for an `einit.ecm` file and load its definitions. Set it to 0 to disable reading this file. Also see the `record-customizations` variable.

`load-fail-ok` System Default: 0

If Epsilon cannot autoload a called EEL function, it will report an error. An EEL subroutine may set this variable nonzero to make Epsilon silently ignore such a function call.

`load-from-state` Default: varies

Epsilon sets this variable to 1 if it loaded its functions from a state file at startup, or 0 if it loaded only from bytecode files.

`locate-path-unix` Preference Default:
"/{*bin*,etc,home*,lib,opt,root,usr*,var*}/**/"

Under Unix, the `locate-file` command uses this variable to decide where to look for a file. It contains part of an extended file pattern that should match those directories where Epsilon should look. The specified file name will be appended to this value. It's a good idea to make sure special file systems like `/proc` are not matched by the pattern.

`mac-framework-dirs` Default: (omitted)

On Macintosh systems, the `find-linked-file` command uses this variable in C mode to help locate header files stored in frameworks. It contains a list of framework directories to search, separated by colons. They may contain wildcard characters.

`macro-runs-immediately` Default: 1

When an EEL function says to run a keyboard macro, Epsilon can do it two ways. Normally Epsilon enters a recursive edit loop, and executes keys from the macro. When the macro ends, Epsilon exits the recursive edit loop, and returns to the EEL function that said to run the macro. This makes keyboard macros behave like functions.

Epsilon can instead simply queue the macro's keys, without employing any loop. Then when an EEL function says to run a keyboard macro, Epsilon just records the fact that it's running a macro and immediately returns to the EEL function. Later when Epsilon is ready for more input (perhaps long after the original macro-queuing function has returned), it begins to use the macro's keys. An EEL function can get this behavior by temporarily setting the `macro-runs-immediately` variable to zero prior to executing a keyboard macro.

`mail-quote-pattern` Preference Default: `"^ [>#] "`

This variable contains a regular expression that matches the start of a line of text quoted email-style. Lines matching this pattern signal the commands `mail-fill-paragraph` and `mail-unquote` that a line is quoted.

`mail-quote-skip` Preference Default: `"^ [>#] +"`

This variable contains a regular expression that matches the portion of a line that `mail-fill-paragraph` should preserve when filling.

`mail-quote-text` Preference Default: `"> "`

This variable contains the text the `mail-quote-region` command inserts at the start of each line to quote it email-style.

`major-mode` System Buffer-specific Default: `"Fundamental"`

This variable holds the name of the current major mode for this buffer.

`margin-right` Preference Buffer-specific Default: 70

This variable holds the current fill column, or right margin. (Also see `c-fill-column` for the C mode equivalent.)

`mark` Buffer-specific Default: none

This variable holds the buffer position of the mark. Several commands operate on the current region. The text between the mark and point specifies the region.

`mark-rectangle-expands` Preference Default: 0

Normally the `mark-rectangle` command begins defining a zero-width rectangle, setting point and mark the same. If this variable is nonzero, that command makes the new rectangle have a width of 1 at the start, by adjusting the current position.

`mark-to-column` Window-specific Default: -1

The window-specific `mark-to-column` variable lets you position the mark in a part of a window where there are no characters. Epsilon uses this variable when it displays a region that runs to the mark's position. It's normally -1, so Epsilon highlights up to the actual buffer position of the mark. If it's non-negative in the current window, Epsilon highlights up to the specified column instead. Epsilon resets `mark-to-column` to -1 whenever the buffer changes, or the mark moves from where it was when you last set `mark-to-column`. (Epsilon only checks these conditions when it redisplay the window, so you can safely move the mark temporarily.)

`mark-unhighlights` Preference Default: 0

When Epsilon is already displaying a highlighted region, region-marking commands like `mark-rectangle` normally change the type of the region. For example, `mark-rectangle` will change a highlighted line region into a rectangular region. If this variable is nonzero, such commands will instead remove the highlighting when Epsilon is already displaying a highlighted region of the desired type. For example, `mark-line-region` will turn off highlighting if Epsilon is displaying a line region. If this variable is zero, Epsilon does nothing when the correct type of highlighted region is already being displayed.

<code>matchdelim</code>	Preference Buffer-specific	Default: 1
If nonzero, typing <code>)</code> , <code>]</code> , or <code>}</code> in C mode displays the corresponding <code>(</code> , <code>[</code> , or <code>{</code> using the <code>show-matching-delimiter</code> command.		
<code>matchend</code>		Default: none
Most of Epsilon's searching primitives set this variable to the far end of the match from the original buffer position.		
<code>matchstart</code>		Default: none
Most of Epsilon's searching primitives set this variable to the near end of the match from the original buffer position.		
<code>max-initial-windows</code>	Preference	Default: 3
When you name several files on Epsilon's command line, Epsilon reads all the named files. But it only displays up to this many in separate windows.		
<code>mem-in-use</code>		Default: varies
This variable holds the amount of space Epsilon is now using for miscellaneous storage (not including buffer text).		
<code>mention-delay</code>	Preference	Default: 0
The <code>mention()</code> primitive displays its message only after Epsilon has paused waiting for user input for <code>mention-delay</code> tenths of a second. When Epsilon prompts for another key, it often displays its prompt in this way.		
<code>menu-bar-flashes</code>	Preference	Default: 2
When you select an item on the menu bar, Epsilon flashes the selected item. This variable holds the number of flashes. (Console versions only.)		
<code>menu-bindings</code>	Preference	Default: 1
If nonzero, Epsilon puts the key bindings of commands into its menu bar. (Console versions only.)		
<code>menu-command</code>	System	Default: varies
When the user selects an item from a menu or the tool bar, Epsilon for Windows returns a special key code, <code>WIN_MENU_SELECT</code> , and sets the <code>menu_command</code> variable to the name of the selected command. Use assignment, not <code>strcpy()</code> or similar, to set this variable.		
<code>menu-file</code>	Preference	Default: "epsilon.mnu"
Epsilon stores the name of the file it is using to display the menu bar in this variable, in all environments except Windows. Also see <code>gui-menu-file</code> .		

<code>menu-stays-after-click</code>	Preference	Default: 1
-------------------------------------	------------	------------

By default, when you click on the menu bar but release the mouse without selecting a command, Epsilon leaves the menu displayed until you click again. Set the `menu-stays-after-click` variable to zero if you want Epsilon to remove the menu when this happens. (Console versions only.)

<code>menu-width</code>	Preference	Default: 35
-------------------------	------------	-------------

This variable contains the width of the pop-up window of matches that Epsilon creates when you press '?' during completion. (Console versions only.)

<code>menu-window</code>	System	Default: none
--------------------------	--------	---------------

This variable stores the window handle of the current menu bar, or zero if there is none. (Console versions only.)

<code>merge-diff-var</code>	Preference	Default: "DIFFVAR"
-----------------------------	------------	--------------------

The `merge-diff` command stores the name of the `#ifdef` variable you select here.

<code>message-history-size</code>	Preference	Default: 1,000,000
-----------------------------------	------------	--------------------

Epsilon keeps a history of prior messages displayed in the echo area in the buffer `#messages#`. The oldest messages are deleted from the top of the buffer whenever it exceeds this size in bytes. If this variable is zero, most commands avoid writing their messages to a `#messages#` buffer.

<code>minimal-coloring</code>	Preference	Default: 0
-------------------------------	------------	------------

Set the `minimal-coloring` variable to 1 to tell Epsilon to limit the amount of coloring it does in order to make code coloring faster. For C files, Epsilon will color all identifiers, keywords, numbers, function calls and punctuation the same, using the `c-ident` color class for all. Epsilon will still uniquely color comments, preprocessor lines, and strings.

<code>misc-language-fill-mode</code>	Preference	Default: 0xffffffff
--------------------------------------	------------	---------------------

Bits in this variable control whether Epsilon will break long comment lines in various language modes. Epsilon will break long comment lines in a language if its corresponding bit is on. The languages and their bits: PostScript 0x1, Conf 0x2, Batch 0x4, Makefile 0x8, Python 0x10, Visual Basic 0x20, VHDL 0x40, Tcl 0x80, Ini files 0x100, NSIS 0x200, and R 0x400.

<code>mode-extra</code>	System	Buffer-specific	Default: none
-------------------------	--------	-----------------	---------------

Epsilon normally includes this text in each mode line of a window displaying this buffer. Internet commands use this to display transfer status via the `set_mode_message()` subroutine. The `mode-more-extra` variable is a non-buffer-specific equivalent.

mode-format

Preference

Default:

`" %b%E%M %l,%c %d%p %s%f "`

This variable specifies what information Epsilon puts on the mode line. Epsilon substitutes values for certain %-sequences in the mode line.

Here are the available sequences:

- %b** The buffer's file name, or if none, the buffer name. If Epsilon displays the buffer name, it usually appears in parentheses. The name may be abbreviated if there's not enough room. The file name will normally be relative to the current directory, but see the `full-path-on-mode-line` variable.
- %B** The buffer's file name as an absolute pathname, or if none, the buffer name.
- %c** The current column number, counting columns from 0.
- %C** The current column number, counting columns from 1.
- %d** The current display column, with a < before it, and a space after. However, if the display column has a value of 0 (meaning horizontal scrolling is enabled, but the window has not been scrolled), or -1 (meaning the window wraps long lines), Epsilon substitutes nothing.
- %D** The current display column, but if the display column is -1, Epsilon substitutes nothing.
- %E** Any "extra text" recorded for this buffer by a command, specified by the `mode-extra` variable, plus any text specified by the `mode-more-extra` variable. These may indicate pending transfer status or similar.
- %f** The name of the current function (in buffers where Epsilon can determine this), or in some modes, other similar information. See the variables `display-definition` and `html-display-definition`.
- %l** The current line number. (Also see the `draw-line-numbers` variable to display line numbers to the left of text.)
- %L** The last component of the buffer's file name (the part after the last / or \), or if none, the last component of the buffer name.
- %m** Epsilon substitutes the text " More ", but only if characters exist past the end of the window. If the last character in the buffer appears in the window, Epsilon substitutes nothing.
- %M** The buffer's mode, inside square brackets.
- %o** The value of point (a character count from the start of the buffer).
- %P** Epsilon substitutes the percentage of point through the buffer, followed by a percent sign.
- %p** Epsilon substitutes the percentage of point through the buffer, followed by a percent sign. However, if the bottom of the buffer appears in the window, Epsilon displays Bot instead (or End if point is at the very end of the buffer). Epsilon displays Top if the top of the buffer appears, and All if the entire buffer is visible.
- %s** Epsilon substitutes "*" if the buffer's modified flag has a nonzero value, otherwise nothing.
- %S** Epsilon substitutes "*" if the buffer's modified flag has a nonzero value, otherwise nothing.
- %h** Epsilon substitutes the current hour in the range 1 to 12.
- %H** Epsilon substitutes the current hour in military time in the range 0 to 23.

%n Epsilon substitutes the current minute in the range 0 to 59.

%e Epsilon substitutes the current second in the range 0 to 59.

%a Epsilon substitutes “am” or “pm” as appropriate.

Note: For the current time, use a sequence like `%2h:%02n%a` for “3:45 pm” or `%02H:%02n:%02e` for “15:45:21”.

%% Epsilon substitutes a literal “%” character.

%< Indicates that redisplay may omit text to the left, if all of the information will not fit.

%> Puts any following text as far to the right as possible.

For any numeric substitution, you may include a number between the % and the letter code, giving the field width: the minimum number of characters to print. You can use the same kinds of field width specifiers as C’s `printf()` function. The sequence `%4c` expands to “`<Space><Space><Space>9`”, `%04c` expands to “`0009`”, and `%-4c` expands to “`9<Space><Space><Space>`”.

Also see the variable `show-when-idle`.

<code>mode-line-at-top</code>	Preference	Default: 0
-------------------------------	------------	------------

If nonzero, Epsilon puts each window’s mode line above the corresponding buffer text.

<code>mode-line-position</code>	Preference	Default: 3
---------------------------------	------------	------------

The `mode-line-position` variable specifies how to position the title text in a tiled window. To set it in an EEL function, use one of the following macros defined in `codes.h`. (These are the same as those used by the `window_title()` primitive.) The `TITLELEFT(n)` macro, which is defined as `(1 + (n))`, positions the title `n` characters from the left edge of the window. The `TITLERIGHT(n)` macro, defined as `-(1 + (n))`, positions the title `n` characters from the right edge of the window. The `TITLECENTER` macro, defined as `0`, centers the title in the window.

<code>mode-more-extra</code>	System	Default: none
------------------------------	--------	---------------

Epsilon normally includes this text in each mode line. Some compilation commands use it to display their status. The `mode-extra` variable is a buffer-specific equivalent.

<code>modified</code>	Buffer-specific	Default: 0
-----------------------	-----------------	------------

If nonzero, the buffer has been modified since it was last read or written.

<code>monochrome</code>	Default: varies
-------------------------	-----------------

This variable is nonzero if Epsilon believes it is running on a monochrome display.

<code>mouse-auto-off</code>	Preference	Default: 1
-----------------------------	------------	------------

If nonzero, Epsilon automatically stops displaying the mouse cursor when you start to type on the keyboard.

`mouse-auto-on` Preference Default: 1

If nonzero, Epsilon automatically displays the mouse cursor when it detects mouse motion.

`mouse-center-yanks` Preference Default: 0

This variable controls the action of the center mouse button, when that button is bound to the `mouse-center` command as it usually is. A value of 0, the default, means the button pans, by calling the `mouse-pan` command. A value of 1 means the button yanks text from a kill buffer or clipboard by calling `mouse-yank`. A value of 2 makes Epsilon yank under Unix, and pan in other environments. For the preceding values, holding down Shift when clicking always produces yanking instead of panning. A value of 3 makes Epsilon ignore pressing the center button.

`mouse-dbl-selects` System Buffer-specific Default: 0

This variable controls what double-clicking with a mouse button does in a particular buffer. If zero, double-clicking selects words. If nonzero, Epsilon instead runs the command bound to the `(Newline)` key.

`mouse-display` Default: none

If nonzero, Epsilon is now displaying a mouse cursor. When Epsilon turns on the mouse cursor automatically due to mouse movement, it sets this variable to 2.

`mouse-mask` Default: 0x2AB (Down or Move events)

Epsilon will only notice and return mouse events that match bits in this variable.

```
0x01 MASK_MOVE
0x02 MASK_LEFT_DN
0x04 MASK_LEFT_UP
0x08 MASK_RIGHT_DN
0x10 MASK_RIGHT_UP
0x20 MASK_CENTER_DN
0x40 MASK_CENTER_UP
0x80 MASK_B4_DN
0x100 MASK_B4_UP
0x200 MASK_B5_DN
0x400 MASK_B5_UP
```

`mouse-panning` System Default: 0

Epsilon uses this variable to help it autoscroll when you click the middle mouse button (on three-button or wheeled mice).

`mouse-pixel-x` Default: none

This variable contains the horizontal mouse position, in the most accurate form Epsilon provides.

`mouse-pixel-y` Default: none

This variable contains the vertical mouse position, in the most accurate form Epsilon provides.

`mouse-screen` System Default: varies

All keys that represent mouse movements or button activity set the `mouse_screen` variable to indicate which screen their coordinates refer to. All tiled windows are on the main screen, screen 0. When Epsilon for Windows creates a dialog box containing one or more Epsilon windows, each Epsilon window has its own screen number.

`mouse-selection-copies` Preference Default: 0

Set this variable to make selecting text with the mouse automatically copy it to a kill buffer, like `copy-region` does. A value of 0 means selecting text doesn't copy it. A value of 1 means selecting text copies it too. A value of 2 makes Epsilon copy under Unix, but not in other environments. When Epsilon copies mouse-selected text to a kill buffer under X11, it also sets the text as X11's primary selection.

`mouse-shift` Default: none

Bits in this variable indicate which shift keys were depressed at the time the current mouse event was enqueued.

`mouse-update-during-scroll` Default: 10,000,000

On buffers bigger than this size, Epsilon won't update the current function name shown in the mode line (see the `display-definition` variable) while you drag the scroll bar box with the mouse, but only when you release the mouse button.

`mouse-x` Default: none

This variable contains the vertical mouse position as a line number on the screen (counting from line zero at the top).

`mouse-y` Default: none

This variable contains the horizontal mouse position as a column number on the screen (counting from column zero on the left).

`mshelp2-collection` Preference Default: ""

Under Windows, the `context-help` command can use the MS-Help2 help engine to display help, if it's installed. That help engine uses various help file sources, called collections. By default, Epsilon tries to select the best collection; set this variable to force it to use a specific collection. The collection name appears in the URL the help engine displays for each topic. (This is an older help format, now rarely used.)

`must-build-mode` Buffer-specific Default: 0

Epsilon "precomputes" most of the text of each mode line, so it doesn't have to figure out what to write each time it updates the screen. Setting this variable nonzero warns Epsilon that mode lines must be rebuilt for all windows displaying this buffer. Epsilon resets the variable to zero after every screen update.

`narrow-end` Buffer-specific Default: 0

Epsilon ignores the last `narrow-end` characters of the buffer, neither displaying them nor allowing any other access to them. But Epsilon does include them when it writes the buffer to a file, and counts them in the total size of the buffer.

`narrow-start` Buffer-specific Default: 0

Epsilon ignores the first `narrow-start` characters of the buffer, neither displaying them nor allowing any other access to them. But Epsilon does include them when it writes the buffer to a file, and counts them in the total size of the buffer.

`national-keys-not-alt` Preference Default: 2

When Epsilon for Unix runs as a curses-style terminal program (not an X11 program), it can interpret key codes in the range 128–255 either as national characters (accented characters) or as Alt versions of other characters. Set this variable to 1 for the former interpretation or 0 for the latter one. Any other value makes Epsilon for Linux provide national characters, and Epsilon for FreeBSD provide Alt keys. (This is intended to accommodate the different console settings on the two systems.) If you need to type accented characters in Epsilon for FreeBSD when it runs outside X, set this variable to 1.

`near-pause` Preference Default: 50

The `find-delimiter` and `show-matching-delimiter` commands pause this many hundredths of a second, when they don't have to reposition the screen to a different part of the buffer in order to show the matching delimiter.

`need-rebuild-menu` System Default: 0

Epsilon sets this nonzero to indicate that it must rebuild the contents of its menu bar.

`net-prompt-get-directory` Preference Default: 0xff

In ssh and telnet buffers, Epsilon tries to set the default directory by parsing the prompt. Bits in this variable control how this happens. All bits are normally enabled.

The 1 bit enables trying to interpret shell prompts.

The 2 bit enables using the `net-prompt-pattern` variable. Without this bit, Epsilon will only look for the special hidden sequence enabled by running one of its shell scripts like `epsilon.sh` on the remote system. See page 134).

The 4 bit controls what happens when Epsilon can't interpret a prompt to get a directory. This bit makes it retain the last directory it received from the process, if any, until it can retrieve another directory. Without this bit, Epsilon will set the buffer's default directory to a URL with no path, whenever it can't parse the most recent output from the host.

The 8 bit makes Epsilon construct an `scp://` URL. Without this bit, it constructs an `ftp://` URL.

The 16 bit says that whenever Epsilon parses a prompt, it should color it using the `process-prompt` color class. Similarly, the 32 bit helps it color user input with the `process-input` color class. Without these bits, Epsilon will use less reliable methods to determine which program output is a prompt.

You can create a host-specific version of this variable by defining a new variable with a name like `net-prompt-get-directory-example-com`. Then Epsilon will use its settings instead for ssh and telnet buffers accessing `example.com`. The 1 bit in `net-prompt-get-directory` itself may be turned off to disable this feature for all hosts.

`net-prompt-pattern` Preference Default: (Omitted)

Epsilon uses this regular expression pattern to match a shell prompt in an ssh or telnet buffer and extract the current directory named by it (if any). The pattern must have its first parenthesized grouping correspond to the directory portion of the prompt.

You can create a host-specific version of this variable by defining a new variable with a name like `net-prompt-pattern-example-com`. See page 134 for more.

`new-buffer-translation-type` Preference Default: 5

When you create a new buffer or file, Epsilon sets its `translation-type` variable to this variable's value. The translation type determines how Epsilon writes or reads a buffer.

A value of 0 (`FILETYPE_BINARY`) makes Epsilon do no line translation, 1 (`FILETYPE_MSDOS`) makes Epsilon strip `<Return>` characters when reading and insert them when writing, 2 (`FILETYPE_UNIX`) makes Epsilon do no line translation, but indicates that the file contains text, 3 (`FILETYPE_MAC`) makes Epsilon replace `<Return>` characters with `<Newline>` characters when reading, and replace `<Newline>` characters with `<Return>` characters when writing.

The default, 5 (`FILETYPE_AUTO`), makes Epsilon use the usual type for this operating system: Unix files under Unix, MS-DOS files elsewhere.

For remote files, those that use a URL syntax, Epsilon uses the `fallback-remote-translation-type` variable instead. Also see the `default-translation-type` variable.

`new-c-comments` Preference Buffer-specific Default: 1

If nonzero, Epsilon creates a comment in C mode using the `//` syntax, rather than the `/* */` syntax. Changing this setting won't affect buffers already in C mode; restarting Epsilon is one way to make the change take effect.

`new-file-ext` Preference Default: ".c"

The `new-file` command creates new buffers with an associated file name that uses this extension. Some modes look at the extension of a buffer's file name to determine how to behave; for example, C mode's syntax highlighting sets its list of keywords differently for C++ buffers than for C buffers.

`new-file-mode` Preference Default: "c-mode"

The `new-file` command creates new buffers set to use this mode. The specified mode-setting command will be run to initialize the buffer.

`new-search-delay` Preference Default: 250

In commands that present a list of choices and automatically search through the list when you type text, Epsilon uses this variable to determine how long a delay must transpire between keystrokes to signal the start of new search text. The delay is in .01 second units. For example, if you type "c", then immediately "o", Epsilon will move to the first entry in the list that starts with "co". But if you pause for more than `new-search-delay` before typing "o", Epsilon begins a new search string and goes to the first entry that starts with "o".

Currently only the `edit-variables` command does this kind of searching.

`normal-cursor` Preference Default: 98099

This variable holds the shape of the cursor in insert mode (as opposed to overwrite mode). It contains a code that specifies the top and bottom edges of the cursor, such as 3006, which specifies a cursor that begins on scan line 3 and extends to scan line 6 on a character box. The topmost scan line is scan line 0.

Scan lines above 50 in a cursor shape code are interpreted differently. A scan line number of 99 indicates the highest-numbered valid scan line (just below the character), 98 indicates the line above that, and so forth. For example, a cursor shape like 1098 produces a cursor that extends from scan line 1 to the next-to-last scan line, one scan line smaller at top and bottom than a full block cursor.

See `normal-gui-cursor` for the Windows or X11 equivalent.

`normal-gui-cursor` Preference Default: 100002

This variable holds the shape of the caret (the text cursor) in insert mode (as opposed to overwrite mode) in the Windows and X11 versions of Epsilon. It contains a code that specifies the height and width of the caret and a vertical offset, each expressed as a percentage of the size of a character in pixels. For example, a width of 100 indicates a caret just as wide as a character. Values close to 0 or 100 are absolute pixel counts, so a width of 98 is two pixels smaller than a character. A width of exactly zero means use the default width.

All measurements are from the top left corner of the character cell. A nonzero vertical offset moves the caret down from its usual starting point at the top left corner.

In EEL programs, you can use the `GUI_CURSOR_SHAPE()` macro to combine the three values into the appropriate code; it simply multiplies the height by 1000 and the offset by 1,000,000, and adds both to the width. So the default Windows caret shape of `GUI_CURSOR_SHAPE(100, 2, 0)`, which specifies a height of 100% of the character size and a width of 2 pixels, is encoded as the value 100,002. The value 100100 provides a block cursor, while 99,002,100 makes a good underline cursor. (It specifies a width of 100%, a height of 2 pixels, and an offset of 99 putting the caret down near the bottom of the character cell.) The `CURSOR_SHAPE()` macro serves a similar purpose for console versions of Epsilon.

The X11 version of Epsilon can only change the cursor shape if you've provided an `Epsilon.cursorstyle:1` resource (see page 6).

`nsis-auto-show-delim-chars` Default: "[()]"

This variable holds the set of delimiter characters that should trigger Epsilon's auto-show-delimiters feature in NSIS mode. Epsilon will search for and highlight the match of each delimiter.

`only-file-extensions` System Default: none

If non-null, file name completion only finds files with extensions from this list. Each extension must include the `.` character and be surrounded by `|` characters.

`opsys` Default: varies

The `opsys` variable tells which operating system version of Epsilon is running, using the following macros defined in `codes.h`. `OS_WINDOWS` (or synonym `OS_DOS`), defined as 1, indicates one of the Windows versions (or else the DOS version) is running. (See the `is-gui` variable to distinguish these.) `OS_OS2`, defined as 2, indicates the OS/2 version is running. `OS_UNIX`, defined as 3, indicates Epsilon is running on a platform with UNIX-derived internals: macOS, Linux, or FreeBSD. (See the `is-unix` variable to distinguish these.)

`over-mode` Preference Buffer-specific Default: 0

If nonzero, typing ordinary characters doesn't insert them between existing characters, but overwrites the existing characters on the line.

`overwrite-cursor` Preference Default: 0099

This variable holds the shape of the cursor in overwrite mode (as opposed to insert mode). See the description of `normal-cursor` for details. See `overwrite-gui-cursor` for the Windows or X11 equivalent.

`overwrite-gui-cursor` Preference Default: 100100

This variable holds the shape of the caret (the text cursor) in overwrite mode (as opposed to insert mode) in the Windows or X11 versions of Epsilon. See the description of `normal-gui-cursor` for details.

`paging-centers-window` Preference Default: 1

If the `paging-centers-window` variable is nonzero, the `next-page` and `previous-page` commands will leave point on the center line of the window when you move from one page to the next. Set this variable to zero if you want Epsilon to try to keep point on the same screen line as it pages. When `paging-centers-window` is zero, these commands won't position point at the start (end) of the buffer when you page up (down) from the first (last) screenful of the buffer, as they normally do.

`paging-retains-view` Default: 0

If the `paging-retains-view` variable is nonzero when Epsilon displays a buffer in a pop-up window, scrolling up or down past the end of the buffer won't remove the pop-up window. Epsilon will ignore attempts to scroll too far.

`path-list-char` Preference Default: ';', or ':' in Unix

This variable contains the character separating the directory names in a configuration variable like `EPSPATH`.

`path-sep` Preference Default: '\', or '/' in Unix

This variable contains the character for separating directory names. It is set to '\' under Windows, '/' under Unix.

`perl-align-contin-lines` Preference Default: 48

By default, the Perl indenter tries to align continuation lines under parentheses and other syntactic items on prior lines. If Epsilon can't find anything on prior lines to align with, or if aligning the continuation line would make it start past column `perl-align-contin-lines`, Epsilon uses a fixed indentation: two levels more than the original line, plus the value of the variable `perl-contin-offset` (normally zero).

Set this variable to zero if you don't want Epsilon to ever try to align continuation lines under syntactic features in previous lines. If zero, Epsilon indents continuation lines by one level (normally one tab stop), plus the value of the variable `perl-contin-offset` (which may be negative).

`perl-auto-show-delim-chars` Preference Default: "[()]"

This variable holds the set of delimiter characters that should trigger Epsilon's auto-show-delimiters feature in Perl mode. Epsilon will search for and highlight the match of each delimiter.

`perl-brace-offset` Preference Default: 0

In Perl mode, Epsilon offsets the indentation of a left brace on its own line by the value of this variable. The `perl-closeback` variable also helps to control this placement.

`perl-closeback` Preference Default: 1

If nonzero, Perl mode aligns a right brace character that ends a block with the line containing the matching left brace character. If zero, Perl mode aligns the right brace character with the first statement inside the block. Additionally, the value 2 makes Epsilon indent lines starting with a right parenthesis so they match the start of the line with its matching left parenthesis, not the left parenthesis itself.

`perl-contin-offset` Preference Default: 0

In Perl mode, Epsilon offsets its usual indentation of continuation lines by the value of this variable. The variable only affects lines that Epsilon can't line up under the text of previous lines.

`perl-detect-expression-pattern` Preference Default: (omitted)

Perl mode coloring uses the regular expression pattern in this variable to help it resolve certain syntax issues with the `/` operator. It matches keywords that typically precede patterns. Perl mode assumes an undecorated identifier not in this list is a subroutine returning a number, not a subroutine taking a pattern expression as an argument.

`perl-heredoc-syntax` Preference Default: 0x200

This variable controls how Epsilon syntax highlights a here document in Perl code. If the heredoc's terminating string ends in `“.html”` or `“.tex”`, Epsilon always colors it using the corresponding mode. Otherwise, it checks this variable. The bit 0x100 makes it try to autodetect TeX code, and the bit 0x200 makes it try to autodetect HTML code. If neither is detected, the bit 1 makes it default to TeX, and 2 makes it default to HTML. (Only include one of these.) If neither bit is present, the here document is colored with the `perl-constant` color class.

`perl-indent` Preference Buffer-specific Default: 0

Perl mode indents each additional level of nesting by this many columns. If the variable is less than or equal to zero, Epsilon uses the value of `tab-size` instead. Set this variable if you want Epsilon to use one number for displaying tab characters, and a different number for indenting Perl code. (Epsilon will indent using a combination of spaces and tabs, as necessary.)

`perl-label-indent` Preference Default: 0

This variable provides the indentation of lines starting with labels in Perl mode. Normally, Epsilon moves labels to the left margin.

`perl-reindent-previous-line` Preference Default: 0

This variable controls whether Epsilon reindents the previous line when you press `<Enter>` in Perl mode.

`perl-tab-override` Preference Default: 8

If you want the width of a tab character in Perl mode buffers to be different than in other buffers, set this variable to the desired value. Perl mode will change the buffer's tab size to the specified number of columns. Setting this variable doesn't change existing buffers; set the `tab-size` variable for that.

`perl-top-braces` Preference Default: 0

Epsilon indents the braces of the top-level block of a function by the number of characters specified by this variable. By default, Epsilon puts such braces at the left margin.

`perl-top-contin` Preference Default: 3

Epsilon indents continuation lines outside of any function body by the number of characters specified by this variable, whenever it cannot find any text on previous lines to align the continuation line beneath.

`perl-top-struct` Preference Default: 8

When a top-level definition appears over several lines, Epsilon indents the later lines by the number of characters specified in this variable, rather than the value of `perl-top-contin`.

`perl-topindent` Preference Default: 1

If nonzero, Epsilon indents top-level statements in a function. If zero, Epsilon keeps such statements at the left margin.

`perldoc-command` Preference Default: "perldoc"

Epsilon's `perldoc` command runs this external program to retrieve Perl documentation.

`permanent-menu` System Default: 0

This variable records whether you want a permanent menu bar. Set it only with the `toggle-menu-bar` command. (GUI versions of Epsilon don't use this variable.)

`permit-window-keys` System Default: 0

Epsilon only recognizes user attempts to scroll by clicking on the scroll bar, or to resize the window, when it waits for the first key in a recursive edit level. Within an EEL command, when an EEL command requests a key, Epsilon normally ignores attempts to scroll, and postpones acting on resize attempts. An EEL command can set the `permit_window_keys` variable to allow these things to happen immediately, and possibly redraw the screen. Bits in the variable control these activities: set the `PERMIT_SCROLL_KEY` bit to permit immediate scrolling, and set `PERMIT_RESIZE_KEY` to permit resizing. Setting the `PERMIT_WHEEL_KEY` bit tells Epsilon to generate a `WIN_WHEEL_KEY` key event after scrolling due to a wheel roll on a Microsoft IntelliMouse.

`php-align-contin-lines` Preference Default: 48

By default, the PHP indenter tries to align continuation lines under parentheses and other syntactic items on prior lines. If Epsilon can't find anything on prior lines to align with, or if aligning the continuation line would make it start past column `php-align-contin-lines`, Epsilon uses a fixed indentation: two levels more than the original line, plus the value of the variable `php-contin-offset` (normally zero).

Set this variable to zero if you don't want Epsilon to ever try to align continuation lines under syntactic features in previous lines. If zero, Epsilon indents continuation lines by one level (normally one tab stop), plus the value of the variable `php-contin-offset` (which may be negative).

`php-auto-show-delim-chars` Preference Default: "[()]"

This variable holds the set of delimiter characters that should trigger Epsilon's auto-show-delimiters feature in PHP mode. Epsilon will search for and highlight the match of each delimiter.

`php-brace-offset` Preference Default: 0

In PHP mode, Epsilon offsets the indentation of a left brace on its own line by the value of this variable. The `php-closeback` variable also helps to control this placement.

`php-closeback` Preference Default: 1

If nonzero, PHP mode aligns a right brace character that ends a block with the line containing the matching left brace character. If zero, PHP mode aligns the right brace character with the first statement inside the block. Additionally, the value 2 makes Epsilon indent lines starting with a right parenthesis so they match the start of the line with its matching left parenthesis, not the left parenthesis itself.

`php-comment-style` Preference Default: 1

When you use the `indent-for-comment` command to insert a comment in PHP code, this variable controls which type of comment Epsilon inserts. The value 1 inserts `#`, the value 2 inserts `//`, and any other value inserts `/*`.

`php-contin-offset` Preference Default: 0

In PHP mode, Epsilon offsets its usual indentation of continuation lines by the value of this variable. The variable only affects lines that Epsilon can't line up under the text of previous lines.

`php-heredoc-syntax` Preference Default: 0x200

This variable controls how Epsilon syntax highlights a here document in PHP code. If the heredoc's terminating string ends in `".html"` or `".tex"`, Epsilon always colors it using the corresponding mode. Otherwise, it checks this variable. The bit 0x100 makes it try to autodetect TeX code, and the bit 0x200 makes it try to autodetect HTML code. If neither is detected, the bit 1 makes it default to TeX, and 2 makes it default to HTML. (Only include one of these.) If neither bit is present, the here document is colored with the `perl-constant` color class. (PHP uses the same color classes as Perl mode.)

`php-indent` Preference Buffer-specific Default: 0

PHP mode indents each additional level of nesting by this many columns. If the variable is less than or equal to zero, Epsilon uses the value of `tab-size` instead. Set this variable if you want Epsilon to use one number for displaying tab characters, and a different number for indenting PHP code. (Epsilon will indent using a combination of spaces and tabs, as necessary.)

`php-label-indent` Preference Default: 0

This variable provides the indentation of lines starting with labels in PHP mode. Normally, Epsilon moves labels to the left margin.

`php-reindent-previous-line` Preference Default: 1

This variable controls whether Epsilon reindents the previous line when you press `(Enter)` in PHP mode.

`php-tab-override` Preference Default: 8

If you want the width of a tab character in PHP mode buffers to be different than in other buffers, set this variable to the desired value. PHP mode will change the buffer's tab size to the specified number of columns. Setting this variable doesn't change existing buffers; set the `tab-size` variable for that.

`php-top-braces` Preference Default: 0

Epsilon indents the braces of the top-level block of a function by the number of characters specified by this variable. By default, Epsilon puts such braces at the left margin.

`php-top-contin` Preference Default: 3

Epsilon indents continuation lines outside of any function body by the number of characters specified by this variable, whenever it cannot find any text on previous lines to align the continuation line beneath.

`php-top-level-indent` Preference Buffer-specific Default: 0

Top-level PHP code will be indented to this column.

`php-top-struct` Preference Default: 8

When a top-level definition appears over several lines, Epsilon indents the later lines by the number of characters specified in this variable, rather than the value of `php-top-contin`.

`php-topindent` Preference Default: 1

If nonzero, Epsilon indents top-level statements in a function. If zero, Epsilon keeps such statements at the left margin.

`point` Buffer-specific Default: none

This variable stores the current editing position. Its value denotes the number of characters from the beginning of the buffer to the spot at which insertions happen.

`position-window-on-screen-line` System Buffer-specific Default: 50

When Epsilon displays a window and discovers that some command has moved `point` to a part of the buffer outside the window, it centers the window around `point`'s new position. Set this variable to change this positioning. It represents the approximate percentage of window lines that should appear above `point`. For instance, a setting of 25 on a 40 line window positions `point` near the window's tenth line.

`postscript-auto-show-delim-chars` Default: "[()]"

This variable holds the set of delimiter characters that should trigger Epsilon's `auto-show-delimiters` feature in PostScript mode. Epsilon will search for and highlight the match of each delimiter.

`preserve-filename-case` Preference Default: 0

Bits in this variable control whether Epsilon displays certain uppercase file names as lowercase, whether uppercase letters in file names are carried over to their corresponding buffer names, and various related settings.

Set the 1 bit to tell Epsilon to use the case of file names exactly as retrieved from the operating system. By default, on case-insensitive file systems like VFAT or NTFS, Epsilon changes all-uppercase file names like `WIN.INI` to lowercase like `win.ini`. It never alters the case of a file name that contains any lowercase letters. (Epsilon never changes the case of filenames under Unix regardless of this bit.)

By default, the buffer name Epsilon chooses when you read a file uses only lowercase letters, even if the corresponding file name uses uppercase. Set the 2 bit if you want buffer names to match the case of their file names. (Setting this variable doesn't rename existing buffers.)

Set the 4 bit if you want Epsilon to fold case when comparing buffer names during completion and `grep`. This and the following flag are most useful when combined with the previous flag.

Set the 8 bit if you want sorting by buffer or file name in a `bufed` listing to fold case.

Set the 16 bit to make Epsilon for Mac OS treat file names as case-sensitive (instead of the default, case-preserving). Epsilon will consider `FILE` and `File` to be two different files.

Set the 32 bit to make Epsilon for Linux or FreeBSD treat file names as case-preserving (instead of the default, case-sensitive). Epsilon will treat `FILE` and `File` as two names for the same file.

`preserve-session` Preference Default: 6

When this variable is 6, Epsilon writes a session file when it exits, and reads one when it starts. Set it to 2 to save the session every time you exit, but not to restore the session by default. Set it to 4 to restore the session normally (see the `session-always-restore` variable) but not to save the session. The value 0 does neither. (The value 1 does both, like 6, for compatibility with previous versions.)

`print-heading` Preference Default: 15

Epsilon for Windows prints a heading at the top of each page. Set this variable to control what it includes. The value 1 makes Epsilon include the file name, 2 makes Epsilon include a page number, and 4 makes Epsilon include the current date. Combine 1 with 8 and Epsilon will display the file's full path instead of just its base name, abbreviating if needed to fit on one line. You can add these values together; the default value includes all the above bit values.

`print-in-color` Preference Default: 1

By default, Epsilon for Windows will print in color on color printers, and in black & white on non-color printers. You can set the `print-in-color` variable to 0, if you don't want Epsilon to ever print in color, or to 2 if you want Epsilon to attempt to use colors even if the printer doesn't appear to be a color printer. (Some printers will substitute shades of grey.) The value 1 produces color printing only on color printers.

`print-line-numbers` Preference Default: 0

Epsilon for Windows will include line numbers in printed output if this variable is nonzero. (To display line numbers on the screen, see the `draw-line-numbers` variable.)

`print-long-lines-wrap` Preference Default: 1

Epsilon for Windows will truncate long lines in printed output if this variable is zero. Otherwise they will be wrapped to the next line.

`print-tabs` Preference Default: 0

If the `print-tabs` variable is zero, Epsilon will make a copy of any text to be printed and convert tab characters within it to spaces, prior to sending it to the printer. If you want Epsilon to send the text to be printed without converting tabs first, set this variable to one.

`process-coloring-rules` Preference Default: 0x7F

This variable controls how Epsilon applies colors to certain elements in concurrent process buffers. The value 0x1 lets Epsilon apply colors to prompts and user input. The value 0x2 lets Epsilon interpret certain ANSI escape sequences that switch output colors. The value 0x4 lets Epsilon adjust the requested background color from such escape sequences for better readability. The value 0x8 makes Epsilon interpret backspaces to make some progress messages and similar look better. The value 0x10 makes Epsilon remove any remaining ANSI escape sequences. The value 0x20 makes Epsilon remove certain xterm-specific escape sequences. The value 0x40 makes Epsilon interpret Unicode UTF-8 encoding in the output. Add these bit values together to enable multiple rules.

`process-completion-dircmds` Preference Default:
`"</word>(cd|chdir|pushd) +"`

When you press `<Tab>` in a process buffer to perform file name completion, Epsilon looks at the word just before the partial file name. If it matches this pattern, Epsilon only offers directory names as completions.

`process-completion-style` Preference Default: 2

You can press `<Tab>` in the process buffer to have Epsilon finish typing the name of a command or file. Bits in this variable control how Epsilon provides that completion when there are multiple matches.

A value of 1 makes Epsilon use Windows-style completion, as in `cmd.exe`; press `<Tab>` to show the first match; press it again to show the next match, or press `Shift-<Tab>` to show the previous match. Without this bit, Epsilon uses Unix-style completion: if `<Tab>` can't add any characters, it tells how many matches there are, and a second press shows all the matches.

A value of 2 makes Epsilon skip the directory entries `.` and `..` and any other entries starting with a `.` when listing matches. This is the default.

A value of 4 modifies how Epsilon does Unix-style completion. By default, if there are a small number of matches, Epsilon displays them in the echo area. If there are many matches, Epsilon lists them in the buffer. With this bit flag, Epsilon always lists the matches in the buffer.

A value of 8 makes Epsilon display the full path of the completed file name when there is only one match.

A value of 16 keeps Epsilon from completing on the names of environment variables, if you begin typing one with a `%` (Windows) or `$` (elsewhere) before it. Epsilon uses its own environment for the list of names, so it won't be able to complete on any names defined only in the shell.

`process-completion-windows-programs` Preference Default: "exe,bat,cmd,com"

Under Windows, Epsilon looks for executable programs on the path to provide completion in a process buffer. It uses this list of extensions to determine which programs are executable. If your command processor is configured to directly run additional types of files, you can modify this list.

`process-current-directory` System Default: varies

When you use a concurrent process, Epsilon stores its current directory in this variable. Setting this variable has no effect.

See the variable `use-process-current-directory` for more details.

`process-echo` Preference Buffer-specific Default: 2

This variable controls how Epsilon handles character echoing in process buffers. Sometimes changing this setting produces better results with particular replacement command processors (shells). A value of 0 makes Epsilon suppress echoing, assuming the subprocess will do it. A value of 1 forces it on, 2 makes Epsilon select an appropriate behavior in a static fashion, and 3 makes Epsilon select an appropriate behavior dynamically, each time it sends input to a process. (Under Unix, 2 and 3 act the same as 1.)

Adding 4 to one of the above values changes which line termination character Epsilon sends, from its usual choice of `<Return>` or `<Newline>` (depending on the platform) to the other one.

Adding 8 makes the filter-region command run the process synchronously; by default it runs synchronously only when replacing text in a buffer, not when sending output to a new buffer.

`process-enter-whole-line` Preference Default: 1

Bits in this variable control what happens when you press `<Enter>` in a process or ssh buffer. The 1 bit makes pressing `<Enter>` in the concurrent process buffer move to the end of the current line before sending it to the process, but only when in a line that has not yet been sent to the process. When you press `<Enter>` from some earlier line in the buffer, the 2 bit makes Epsilon copy the entire current line (except for the prompt) to the end of the buffer (making it easier to repeat a command), instead of inserting a new line. The 4 bit makes Epsilon position the cursor at the start of a line retrieved from command history with `Alt-n` or `Alt-p`, not at its end. (This last bit also affects telnet buffers.)

`process-exit-status` System Buffer-specific Default: varies

Epsilon sets this variable when a concurrent process exits, to indicate its exit code. Before the process exits, it contains the value `PROC_STATUS_RUNNING`.

`process-max-size` Preference Buffer-specific Default: -1

If this variable contains a positive number, Epsilon will delete lines from the start of a process buffer whose size exceeds this value.

`process-next-error-options` Preference Default: 7

Commands like `next-error` parse compiler error messages and load the file they indicate. They try to determine the directory of the file based on the output of the compiler, plus directory-switching output from building tools like `Make`, but sometimes they may not be able to determine the right directory. Bits in this variable say what to do when that happens.

Bit 1 enables finding an already-loaded file with the same basename, regardless of its directory. Bit 2 makes Epsilon display a message whenever this happens. If Epsilon still can't find the file, bit 4 makes it display a message instead of creating a new buffer for the file.

`process-output-to-window-bottom` Preference Default: 1

When output arrives from the concurrent process, Epsilon scrolls the text so the end of the output appears on the last line of the window, if possible, like a traditional console window. Set this variable to zero to disable this feature.

`process-pass-drive-directories` Preference Default: 0

When Epsilon for Windows starts a process, it passes its own current directory settings to the process. Each drive has its own current directory setting. By default, Epsilon doesn't pass its current directory setting for any network or removable drives, because passing current directory information for these types of drives can be slow. Set this variable to 1 if you want Epsilon to pass its current directory setting for each network drive, 2 for each removable drive, or 3 for both.

`process-prompt-pattern` Preference Default: ""

Under Windows NT/2000/XP and later Windows versions, certain process buffer functions like completion, or interpreting compiler error messages, require Epsilon to determine the current directory of the command processor running there. Epsilon does this by examining each prompt

from `cmd.exe`, such as `C:\WINNT>`. If you've set a different format for the prompt, set this variable to tell Epsilon how to retrieve the directory name from it.

If this variable is nonempty, it must be a regular expression pattern that matches a prompt. The first pair of parentheses (not counting any using the syntax “(?:)”) must match the directory portion of the prompt. Epsilon will try to parse the buffer using this pattern starting from the beginning of the last line of the prompt. The pattern may either match from that position forward, or from that position backward (to handle multi-line prompts).

As an example, say you use Cygwin's bash shell and set the `PS1` prompt variable to `'\u@\h\w\n$ '`, producing a two-line prompt, with a user name, an at sign, a machine name, and the current directory on the first line. Set `process-prompt-pattern` to `@<^wspace>+(.)<newline>`. This uses the machine name to locate the start of the current directory on the previous line. The default `cmd.exe` prompt of `PG` could use a pattern of `“(.*)%>”`, which selects all the text on a line before a final `>` as the directory name.

In non-Windows environments, Epsilon uses a different method to determine the command processor's current directory, so setting this variable is unnecessary.

<code>process-spell-word-pattern</code>	Preference	Default: (omitted)
---	------------	--------------------

The Spell minor mode uses this regular expression pattern to find misspelled words in Process mode. It doesn't exclude any words based on syntax rules.

<code>process-tab-size</code>	Preference	Default: 8
-------------------------------	------------	------------

Epsilon sets the displayed width of `<Tab>` characters in the process buffer to this value.

<code>process-view-error-lines</code>	Preference	Default: 0
---------------------------------------	------------	------------

If this variable is nonzero, it tells the `next-error` and `previous-error` commands to not only show the source line in error, but the process buffer containing the error message. It shows the latter in a window at the bottom of the screen whose height is determined by this variable.

You can force Epsilon to wrap or scroll long lines in such a window by setting the `wrap-view-error-lines` variable.

<code>process-warn-on-exit</code>	Preference	Default: 0
-----------------------------------	------------	------------

If this variable is set to 1, whenever you try to exit Epsilon and a concurrent process is running (or there's a connected telnet or ssh buffer), Epsilon will use the `exit-process` command to try to make it exit. If the process refuses to exit, Epsilon will display an error message, and won't exit. If this variable is set to 2 and there's a running process or a connected telnet or ssh buffer, Epsilon will prompt for permission to exit, but it won't use `exit-process`.

<code>process-warn-on-killing</code>	Preference	Default: 0
--------------------------------------	------------	------------

By default, when you kill a buffer with an active concurrent process (or a telnet or ssh buffer with an active connection), Epsilon asks it to exit using the `exit-process` command, which types “exit” to it. If that doesn't cause it to exit, Epsilon asks if you want to kill the buffer anyway, disconnecting it from the process (which it tries to kill in stronger ways). (With telnet buffers, Epsilon simply disconnects, and doesn't send an “exit” command.)

Set this variable to 1 if you want Epsilon to prompt before doing anything, when a buffer you've said to kill contains an active process. If you confirm that it's OK to try to exit the process, Epsilon will then proceed as above.

Set this variable to 2 if you want Epsilon to simply display an error message when you try to kill a buffer with an active process.

`process-yank-confirm`

Preference

Default: 2

If you yank this number of lines or more at a process buffer prompt, Epsilon will ask if you really mean to have them immediately executed. Set this variable to zero to disable this warning.

`prompt-with-buffer-directory`

Preference

Default: 2

The `prompt-with-buffer-directory` variable controls how Epsilon uses the current directory at file prompts. When this variable is 2, the default, Epsilon inserts the current buffer's directory at many file prompts. This makes it easy to select another file in the same directory. You can edit the directory name, or you can begin typing a new absolute pathname right after the inserted pathname. Epsilon will delete the inserted pathname when it notices your absolute pathname. This behavior is similar to Gnu Emacs's.

A setting of 3 makes Epsilon insert the current buffer's directory in the same way, but keeps Epsilon from automatically deleting the inserted pathname if you type an absolute one.

When `prompt-with-buffer-directory` is 1, Epsilon temporarily changes to the current buffer's directory while prompting for a file name, and interprets file names relative to the current directory. This behavior is similar to the "pathname.e" extension available for previous versions of Epsilon.

When `prompt-with-buffer-directory` is 0, Epsilon doesn't do anything special at file prompts. Relative file names will be interpreted based on the global current directory set with the `cd` command, not any directory associated with the current buffer.

`pull-word-from-tags`

Preference

Default: 2

This variable determines whether the `pull-word` command will look for possible completions in the current tag file. If 0, it won't. If 1, it will, but only if a tag file is already loaded. If 2, it will load and tag file if necessary and then look.

`push-cmd`

Preference

Default: "make"

This variable holds the default command line for running another program. The variable is a template based on the current file name, so you can set it to automatically operate on the current file. Epsilon substitutes pieces of the current file name for codes in the template, as follows (examples are for the file `c:\dos\read.me`):

- %p** The current file's path (`c:\dos\`).
- %b** The base part of the current file name (`read`).
- %e** The extension of the current file name (`.me`).
- %f** The full name of the current file (`c:\dos\read.me`).
- %r** The name of the file relative to the current directory. (`read.me` if the current directory is `c:\dos`, `dos\read.me` if the current directory is `c:\`, otherwise `c:\dos\read.me`).
- %x** The full pathname of the directory containing the Epsilon executable.
- %X** The full pathname of the directory containing the Epsilon executable, after converting all Windows long file names to their equivalent short name aliases.

`push-cmd-unix-interactive` Preference Default: "xterm &"

When Epsilon for Unix runs as an X11 program, the push command executes this command when it needs to start an interactive shell. (See `push-cmd-unix-macos-interactive` for the macOS version.)

`push-cmd-unix-macos-interactive` Preference Default:
"xterm || /opt/X11/bin/xterm &"

When Epsilon for Unix runs as an X11 program under macOS, the push command executes this command when it needs to start an interactive shell.

`python-auto-show-delim-chars` Preference Default: "([])"

This variable holds the set of delimiter characters that should trigger Epsilon's auto-show-delimiters feature in Python mode. Epsilon will search for and highlight the match of each delimiter.

`python-comment-region-options` Preference Default: 0

Bits in this variable modify how the comment-region command behaves in Python mode. See the description of the `comment-region-options` variable.

`python-delete-hacking-tabs` Preference Default: 10

This variable controls the behavior of the (Backspace) key in Python mode. It can be set to change tabs to spaces, or to delete multiple spaces at once under certain conditions.

See the `delete-hacking-tabs` variable for details on the meaning of the various bits.

`python-indent` Preference Buffer-specific Default: 4

Each level of indentation in Python mode will occupy this many columns.

`python-indent-to-comment` Preference Default: 1

If nonzero, Python mode typically indents each line to match the previous nonblank line. If zero, Python mode typically indents each line to match the previous nonblank, noncomment line.

`python-indent-with-tabs` Preference Default: 0

If nonzero, Python mode indents using tabs. This variable applies to newly-created buffers only, and serves to initialize the `indent-with-tabs` variable in Python mode. Set the latter to change this setting for the current buffer only.

`python-language-level` Preference Default: 3

Set this variable to 2 if you want syntax highlighting to color only Python 2 keyword names, not those only in Python 3.

`python-tab-override` Preference Default: 8

If you want the width of a tab character in Python buffers to be different than in other buffers, set this variable to the desired value. Python mode will change the buffer's tab size to the specified number of columns. Setting this variable doesn't change existing buffers; set the `tab-size` variable for that.

`quiet-write-state` System Default: 0

If this variable is nonzero, the `write-state` command won't prompt, but simply write the state to the file it would offer as a default.

`readonly-pages` Preference Default: 1

In a read-only buffer you can use the `<Space>` and `<Backspace>` keys to page forward and back more conveniently. Other inserting keys display an error message. Set this variable to zero if you want these keys to display an error message, not page.

`readonly-warning` Preference Default: 3

Bits in this variable control Epsilon's action when it reads a read-only file:

- `ROWARN_MSG (1)` Epsilon displays a warning message.
- `ROWARN_BUF_RO (2)` Epsilon sets the buffer read-only.
- `ROWARN_BELL (4)` Epsilon beeps.
- `ROWARN_GREP (8)` Postpone the above actions during multi-file search.
- `ROWARN_REFRESH (16)` Postpone the above actions during refresh-files.

Add these together to get multiple actions. Epsilon uses the `ROWARN_GREP` and `ROWARN_REFRESH` bits internally, to alter its own behavior during the execution of particular commands. They're not intended for use in other contexts.

`recall-id` System Default: none

Epsilon's line input subroutines let you recall previous responses to each prompt. Epsilon normally keeps track of which responses go with which prompts by recording the type of response (file name, buffer name, etc.) and the name of the command that prompted for the text. A command can tell Epsilon to use a different "handle" for a prompt by setting the `recall-id` variable to a string containing the handle. For example, if you wrote three new EEL commands and wanted them to share previous responses, you could include the line `save_var recall_id = "my_responses";` in each command prior to calling the input function.

`recall-longest-response` Preference Default: 10,000

Epsilon saves your responses to each prompt, except those longer than specified by this variable.

`recall-maximum-session` Preference Default: 500,000

Epsilon saves previous responses to all prompts in its session file, so you don't have to type them in again. It uses up to `recall-maximum-session` characters in a session file for previous responses, discarding the oldest unrecalled responses when necessary.

`recall-maximum-size` Preference Default: 500,000

Epsilon saves previous responses to all prompts, so you don't have to type them in again. It retains up to `recall-maximum-size` bytes of previous responses, discarding the oldest unrecalled responses when necessary.

`recall-prior-response-options` Preference Default: 0

Epsilon remembers your previous responses at virtually all prompts, and can recall them. The `recall-prior-response-options` variable controls how Epsilon does this. Bits in the variable control particular aspects of Epsilon's behavior. Add them together to customize the behavior you want.

By default, keys like `<Up>` at a prompt display a list of all previous responses to that prompt. You can then select one, maneuvering in the list using normal Epsilon commands. Set the 1 bit to make `<Up>` and `<Down>` instead substitute prior responses in place, without displaying a list. You can still use `Alt-<Up>` to display the list.

Set the 2 bit to do the same for search prompts. (This doesn't change Epsilon's behavior during incremental search, only in non-incremental mode, `grep`, `replace` commands, and so forth.)

If you've set Epsilon to substitute prior responses in place, pressing `<Up>` or `<Down>` repeatedly eventually reaches the oldest or youngest response to that prompt. Set the 4 bit if you want these keys to wrap around to the other end of the list when that happens, so selecting older and older responses will eventually reach the youngest again.

If you've set Epsilon to substitute prior responses in place, set the 8 bit to make Epsilon include a blank response at the very end of the list of responses.

At any prompt, `Ctrl-p`, `Ctrl-n` and `Ctrl-Alt-p` are equivalent to `<Up>`, `<Down>`, and `Alt-<Up>`, respectively.

`recall-response-selected` Preference Default: 0

When you use a key like `Alt-e` or `Ctrl-Alt-p` that inserts a prior response at a prompt, it will be highlighted as a selection if this variable is nonzero, so you can replace it by typing any text, or retain it by using some other editing key like `<Right>`.

`recognize-password-pattern` Preference Default: (omitted)

In telnet, ssh and concurrent process buffers, Epsilon looks for a `Password:` prompt and intercepts it to help hide your password. Epsilon uses this regular expression pattern to recognize when text is a password prompt. Also see the `recognize-password-prompt` variable.

`recognize-password-prompt` Preference Buffer-specific Default: 3

In telnet, ssh and concurrent process buffers, Epsilon looks for a `Password:` prompt and intercepts it to help hide your password. Set this variable to zero if you don't want this feature. Set it to 1 if you want it only in telnet buffers, 2 if you want it only in concurrent process buffers, or 3 if you want both. If you disable this feature (or it doesn't recognize an unusual password prompt), you can use the `send-invisible` command to manually send a password without letting it appear in the buffer. Also see the `recognize-password-pattern` variable.

`record-customizations` Preference Default: 0

When you set a variable, change color settings, define a macro, or similar, Epsilon can record the setting in a customizations file named `init.ecm` so it will be available in future sessions. The customizations file normally resides in a buffer named “-customizations”, and Epsilon automatically saves it without prompting before you exit Epsilon, and at other times.

Set this variable to 1 if you want Epsilon to record customization settings in the “-customizations” buffer but not save them automatically; you can save them like any other file, and Epsilon will remind you to do so before exiting if there are unsaved customizations.

Set it to 2 and Epsilon will automatically save your customizations without prompting before you exit Epsilon, and at other times.

With a setting of 0, Epsilon won’t automatically record customization settings in this way. You can use the `list-customizations` command at any time to prepare a list of current customizations, ready to save.

Also see the `load-customizations` variable.

`recording-suspended` System Default: 0

The `pause-macro` command sets this variable nonzero to indicate that it has suspended recording of a keyboard macro.

`regex-first-end` Preference Default: 0

If nonzero, Epsilon’s standard regular expression searching commands find the match of the pattern that ends first, rather than the one that begins first.

`regex-shortest` Preference Default: 0

If nonzero, Epsilon’s standard regular expression searching commands find the shortest match of the pattern, rather than the longest match.

`reindent-after-c-yank` Preference Default: 10000

When you yank text into a buffer in C mode, Epsilon automatically reindents it. This is similar to the “smart paste” feature in some other editors. Epsilon won’t automatically reindent very large blocks of text. This variable specifies the size in characters of the largest block that should automatically be reindented. Set it to 0 to disable automatic reindent in C mode, or -1 to reindent all text yanked in C mode.

Also see the variables `reindent-c-comments` and `reindent-one-line-c-comments`.

`reindent-after-perl-yank` Preference Default: 0

When you yank text into a buffer in Perl mode, Epsilon automatically reindents it. This is similar to the “smart paste” feature in some other editors. Epsilon won’t automatically reindent very large blocks of text. This variable specifies the size in characters of the largest block that should automatically be reindented. Set it to 0 to disable automatic reindent in Perl mode, or -1 to reindent all text yanked in Perl mode.

`reindent-after-vbasic-yank` Preference Default: 0

When you yank text into a buffer in Visual Basic mode, Epsilon automatically reindents it. This is similar to the “smart paste” feature in some other editors. Epsilon won’t automatically reindent very large blocks of text. This variable specifies the size in characters of the largest block that should automatically be reindented. Set it to 0 to disable automatic reindent in Visual Basic mode, or -1 to reindent all text yanked in Visual Basic mode.

`reindent-after-yank` Preference Default: 0

This variable controls whether Epsilon automatically reindents blocks of text you yank into the current buffer. This is similar to the “smart paste” feature in some other editors. This variable specifies the size in characters of the largest block that should automatically be reindented. A value of 0 disables automatic reindent in this buffer, and -1 removes any size limitation. Mode-specific variables like `reindent-after-c-yank` take precedence over this variable.

`reindent-c-comments` Preference Buffer-specific Default: 17

Bits in this variable control how C mode indents lines inside a comment. For block comments, this includes the line that starts the block comment only if its ‘/*’ occurs at the beginning of the line.

The 2 bit makes Epsilon reindent block comment lines when yanking a block of text and automatically reindenting it.

The 4 bit makes Epsilon reindent a block comment line when you press `<Enter>` at the end of one, and `c-reindent-previous-line` would ordinarily permit Epsilon to reindent it. (If auto-fill is on, its rules will prevent such reindenting regardless of this bit.)

The 8 bit makes Epsilon reindent a one-line comment (a line beginning with `//`) when you press `<Enter>` at the end of one and `c-reindent-previous-line` would ordinarily permit Epsilon to reindent it.

The 16 bit lets the indent-region reindent block comment lines.

Finally, the 1 bit permits reindenting block comment lines during any other commands that cause reindenting.

`reindent-c-preprocessor-lines` Preference Buffer-specific Default: 1

Set this variable to zero if you don’t want Epsilon to change the indentation of preprocessor lines when indenting code in C mode. Set it to 2 to indent preprocessor lines to the same indent level that any other line would receive. The default value 1 forces preprocessor lines to the left margin.

`reindent-one-line-c-comments` Preference Buffer-specific Default: 1

This variable controls how Epsilon indents comment lines that start with `/*`. If 0, Epsilon never changes the indentation of these lines in commands like `indent-region`. If 1, Epsilon reindents these lines, except when yanking a block of text and automatically reindenting it. If 2, Epsilon reindents in all cases.

`reindent-perl-comments` Preference Buffer-specific Default: 1

Set this variable to zero to keep commands like `indent-region` from reindenting Perl comment lines.

<code>replace-by-case</code>	Preference	Default: 1
<p>This variable controls how commands like <code>query-replace</code> and <code>regex-replace</code> modify the case of the replacement text to match the case of individual matches.</p> <p>If 0, the replacement text is inserted as-is.</p> <p>If 1, the search is case-insensitive, and neither the search nor the replace text contain capital letters, Epsilon makes the replacement text upper case or capitalized whenever a match is entirely upper case or capitalized, respectively.</p> <p>If 2, the Ctrl-c to toggle case-folding in searches includes a third option, <code>CaseRepl</code>. When this option is toggled on and searching is case-insensitive, Epsilon modifies the case of the replacement text based on the case of the match, according to these rules:</p> <p>If the match contains both uppercase and lowercase characters, the first character of the replacement text will be given the case of the first character of the match. Otherwise, if the match contains only lowercase or only uppercase characters, the replacement text will be made all lowercase or uppercase, respectively. Otherwise, the case of the replacement text will not be altered.</p>		
<code>replace-in-region</code>	Preference	Default: 0
<p>If nonzero, whenever a region is highlighted when you invoke a replacing command like <code>query-replace</code> and <code>regex-replace</code>, its effect will be limited to that region. After replacing, the region will remain highlighted if this variable is set to 2, but not if it's set to 1.</p>		
<code>replace-num-changed</code>	System	Default: 0
<p>The <code>string_replace()</code> subroutine sets the <code>replace-num-changed</code> variable to the number of matches it changed.</p>		
<code>replace-num-found</code>	System	Default: 0
<p>The <code>string_replace()</code> subroutine sets the <code>replace-num-found</code> variable to the number of matches it found.</p>		
<code>resize-menu-list</code>	System	Default: 0
<p>An EEL completion function can set this variable nonzero to indicate that if the user tries to list possible completion choices, the window displaying the choices should be widened if necessary to fit the widest choice. This variable has no effect on Epsilon windows within GUI dialogs.</p>		
<code>resize-rectangle-on-tab</code>	Preference	Default: 0
<p>By default, pressing <code><Tab></code> or <code>Shift-<Tab></code> while defining a rectangular region shifts the region right or left. Set this variable to 1 if you want these keys to resize the region without changing the text, by positioning to the next or previous tab stop.</p>		
<code>restart-concurrent</code>	Preference	Default: 1
<p>When the <code>push</code>, <code>make</code>, or <code>compile-buffer</code> commands exit from a concurrent process to run a command non-concurrently, they will restart the concurrent process once the command finishes. Set <code>restart-concurrent</code> to zero if you don't want Epsilon to restart the concurrent process in this case.</p>		

<code>restore-color-on-exit</code>	Preference	Default: 1
------------------------------------	------------	------------

If nonzero, the Windows Console version of Epsilon tries to restore the screen color when you exit. If zero, Epsilon tries to set the color to the after-exiting color class, as specified with the `set-color` command. (Sometimes the operating system environment overrides this and forces a particular color.)

<code>resynch-match-chars</code>	Preference	Default: 15
----------------------------------	------------	-------------

If you invoke `compare-windows` again immediately after it has found a difference, the command will try to resynchronize the windows by moving forward in each window until it finds a match of at least `resynch-match-chars` characters.

<code>return-raw-buttons</code>	System	Default: 0
---------------------------------	--------	------------

If you click a button in a dialog under Epsilon for Windows, Epsilon represents the input with an ordinary key value, such as `<Enter>` when you click an Ok button. An EEL program can temporarily set this variable to a nonzero value to retrieve button presses as distinct keys. All buttons will then appear with the key code `WIN_BUTTON`. Use the `key-is-button` variable to distinguish one button from another.

<code>rev-search-key</code>	Preference	Default: -1
-----------------------------	------------	-------------

Inside a search command, Epsilon recognizes a key with this key code as a synonym for `Ctrl-r`, for pulling in a default search string or changing the search direction.

<code>rmode-auto-show-delim-chars</code>	Default: "([])"
--	-----------------

This variable holds the set of delimiter characters that should trigger Epsilon's `auto-show-delimiters` feature in R mode. Epsilon will search for and highlight the match of each delimiter.

<code>rmode-indent</code>	Preference Buffer-specific	Default: 4
---------------------------	-------------------------------	------------

Each level of indentation in R mode will occupy this many columns.

<code>rmode-indent-to-comment</code>	Preference	Default: 1
--------------------------------------	------------	------------

If nonzero, R mode typically indents each line to match the previous nonblank line. If zero, R mode typically indents each line to match the previous nonblank, noncomment line.

<code>rmode-tab-override</code>	Preference	Default: 4
---------------------------------	------------	------------

If you want the width of a tab character in R mode buffers to be different than in other buffers, set this variable to the desired value. A value of 0 means use `tab-size`'s default value. R mode will change the buffer's tab size to the specified number of columns. Setting this variable doesn't change existing buffers; set the `tab-size` variable for that.

<code>run-by-mouse</code>	System	Default: 0
---------------------------	--------	------------

If nonzero, this command was run by the mouse, via the menu bar or tool bar.

`save-all-without-asking` Preference Default: 0

Set this variable nonzero if you want the `save-all-buffers` command to skip over those buffers created with the File/New menu item or `new-file` command that still lack associated file names. Instead of prompting for a file name, it will report which buffers it didn't save.

`save-when-making` Preference Default: 2

If zero, the `make` command doesn't warn about unsaved buffers before running another program. If one, the command automatically saves all unsaved buffers without asking. If two, Epsilon asks if you want to save the unsaved buffers.

`scp-client-style` Preference Default: 0

By default, Epsilon runs the `sftp` program to perform all `scp` operations. To access very old `ssh` servers, it can instead use the provided `epsilon-xfer-helper` program over a secure link, if you first set things up. Set this variable to 0x1 to make Epsilon issue commands for the helper program, not `sftp`. With this setting, Epsilon will run a separate `scp` program each it has to copy a file.

If you have trouble viewing `dired` listings, you may have an older version of the `sftp` program; set this variable to 0x2 to have Epsilon use old-style `sftp` commands.

When you type an `scp://` URL that doesn't end in a `/` character, Epsilon first checks to see if it's a directory. You can have Epsilon skip this check by setting the 0x8 bit. With this option, whenever you type a directory name, you must either end it with a `/` character to indicate it's a directory, or use the `dired` command instead of one of the other file-reading commands.

See the main SCP topic for details.

`scp-list-flags` Preference Default: "-la"

When you use an `scp://` URL and Epsilon runs an `sftp` client to list files in a directory, it provides these flags to `sftp`'s `ls` command. If you connect to an old `sftp` server, it may not understand these flags. Use "-l" instead, though listings may then omit hidden files.

`scp-read-file-no-user-template` Preference Default: "scp -q %h:%f %l"

Normally Epsilon copies files using an `sftp` program, but you can set it to use a separate program instead, if you must use an old `ssh` system that lacks `sftp` support. When Epsilon tries to access an `scp` URL, it uses this template to build a command that copies a file from a remote host to a local file, in cases where it can't use an `sftp` program to do this. (See the `scp-client-style` variable.)

Epsilon substitutes the host name for `%h`, the file name part for `%f`, and the name of the local destination file for `%l`. The local file will be a relative pathname to a file in the current directory of the `scp` command. The resulting command line should transfer the file quietly, without progress messages.

If the `scp` URL contains a user name, Epsilon uses the `scp-read-file-template` variable instead. Epsilon checks for host-specific variables before checking this variable.

<code>scp-read-file-template</code>	Preference	Default: <code>"scp -q %u@%h:%f %l"</code>
-------------------------------------	------------	--

Normally Epsilon copies files using an sftp program, but you can set it to use a separate program instead, if you must use an old ssh system that lacks sftp support. When Epsilon tries to access an scp URL, it uses this template to build a command that copies a file from a remote host to a local file, in cases where it can't use an sftp program to do this. (See the `scp-client-style` variable.)

Epsilon substitutes the host name for `%h`, the user name for `%u`, the file name part for `%f`, and the name of the local destination file for `%l`. The local file will be a relative pathname to a file in the current directory of the scp command. The resulting command line should transfer the file quietly, without progress messages.

If the scp URL doesn't include a user name, Epsilon uses the `scp-read-file-no-user-template` variable instead. Epsilon checks for host-specific variables before checking this variable.

<code>scp-run-helper-no-user-template</code>	Preference	Default: <code>"%c %h"</code>
--	------------	-------------------------------

When Epsilon tries to perform an operation on an scp URL, such as listing the contents of the directory to which the scp URL refers, file name completion, or similar, it uses this template to build a command to execute. With default settings, it runs the sftp program, which must already be installed. See the main SCP topic for details.

Epsilon substitutes the host name for `%h`. It substitutes the contents of either the `scp-unix-sftp-command` or `scp-windows-sftp-command` variable for the `%c`, depending on the OS where Epsilon is running.

If the scp URL contains a user name, Epsilon uses the `scp-run-helper-template` variable instead. Epsilon checks for host-specific variables before checking this variable.

<code>scp-run-helper-template</code>	Preference	Default: <code>"%c %u@%h"</code>
--------------------------------------	------------	----------------------------------

When Epsilon tries to perform an operation on an scp URL, such as listing the contents of the directory to which the scp URL refers, file name completion, or similar, it uses this template to build a command to execute. With default settings, it runs the sftp program, which must already be installed. See the main SCP topic for details.

Epsilon substitutes the host name for `%h` and the user name for `%u`. It substitutes the contents of either the `scp-unix-sftp-command` or `scp-windows-sftp-command` variable for the `%c`, depending on the OS where Epsilon is running.

If the scp URL doesn't include a user name, Epsilon uses the `scp-run-helper-no-user-template` variable instead. Epsilon checks for host-specific variables before checking this variable.

<code>scp-unix-sftp-command</code>	Preference	Default: <code>"sftp"</code>
------------------------------------	------------	------------------------------

When Epsilon for Unix constructs a command line to perform some scp operation, it substitutes the value of this variable for any `%c` sequence in the command line template. Any `%x` in the setting will be replaced with the name of the directory containing Epsilon's executable; any `%X` will be replaced by the 8.3 version of that directory name. Also see the `scp-windows-sftp-command` variable.

`scp-windows-sftp-command` Preference Default: "%w sftp"

When Epsilon for Windows constructs a command line to perform some scp operation, it substitutes the value of this variable for any %c sequence in the command line template.

Any %w in the setting will be replaced by an invocation of an appropriate helper program when running on a system with Cygwin installed (or with nothing if Cygwin is not installed). (The helper program instructs any programs that follow on the command line to behave as if running interactively, when sending output or prompting for user input. In particular, this makes it prompt for a password or passphrase.)

Any %x in the setting will be replaced with the name of the directory containing Epsilon's executable; any %X will be replaced by the 8.3 version of that directory name. The default setting runs Cygwin's sftp program using a small helper program that makes it prompt for a passphrase or password in a way that Epsilon can intercept. Also see the `scp-unix-sftp-command` variable.

`scp-write-file-no-user-template` Preference Default: "scp -q %l %h:%f"

Normally Epsilon copies files using an sftp program, but you can set it to use a separate program instead, if you must use an old ssh system that lacks sftp support. Epsilon uses this template to build a command that copies a local file to a remote host using scp, in cases where it can't use an sftp program to do this. (See the `scp-client-style` variable.)

Epsilon substitutes the host name for %h, the destination file name part for %f, and the name of the local file for %l. The local file will be a relative pathname to a file in the current directory of the scp command. The resulting command line should transfer the file quietly, without progress messages.

If the scp URL contains a user name, Epsilon uses the `scp-write-file-template` variable instead. Epsilon checks for host-specific variables before checking this variable.

`scp-write-file-template` Preference Default: "scp -q %l %u@%h:%f"

Normally Epsilon copies files using an sftp program, but you can set it to use a separate program instead, if you must use an old ssh system that lacks sftp support. Epsilon uses this template to build a command that copies a local file to a remote host using scp, in cases where it can't use an sftp program to do this. (See the `scp-client-style` variable.)

Epsilon substitutes the host name for %h, the user name for %u, the destination file name part for %f, and the name of the local file for %l. The local file will be a relative pathname to a file in the current directory of the scp command. The resulting command line should transfer the file quietly, without progress messages.

If the scp URL doesn't include a user name, Epsilon uses the `scp-write-file-no-user-template` variable instead. Epsilon checks for host-specific variables before checking this variable.

`screen-cols` System Default: varies

This variable holds the number of columns on the screen.

`screen-lines` System Default: varies

This variable holds the number of lines on the screen.

`scroll-at-end` Preference Default: 1

When you move past the top or bottom edge of the window via the up-line or down-line commands, Epsilon scrolls the screen by this many lines. If `scroll-at-end` is zero, Epsilon instead centers the new line in the window.

`scroll-bar-type` Preference Default: 1

Epsilon for Windows can display two types of scroll bars. By default `scroll-bar-type` is 1, and Epsilon uses a line-based approach, with a “thumb” size that varies to reflect the number of lines visible in the window relative to the number of lines in the buffer. On extremely large buffers, this could be slow, so you can set the variable to 0 and Epsilon will use a fixed-size thumb as in previous versions.

`scroll-init-delay` Preference Default: 35

Epsilon delays `scroll-init-delay` hundredths of a second after its first scroll due to a mouse click on the scroll bar, before it begins repeatedly scrolling at `scroll-rate` lines per second.

`scroll-rate` Preference Default: 45

Epsilon scrolls by this many lines per second when scrolling due to mouse movements.

`search-defaults-from` Preference Default: 1

This variable controls what initial search string appears when you use a searching command like `incremental-search` or `query-replace`.

If 0, Epsilon never provides an initial search string.

If 1, whenever there’s a region highlighted when you invoke the search command, that becomes the initial search string; in other cases there’s no initial search string. The highlighted region is ignored when it’s very long, a macro is running, the `typing-deletes-highlight` variable is zero, or Epsilon has been set to restrict searching to a highlighted region with `search-in-region` or `replace-in-region`.

If 2, the search string from the previous searching command is always the initial search string.

`search-delete-match` Preference Default: 1

If nonzero, then during incremental searching, keys bound to the `delete-character` command, like Del or Ctrl-d, will delete the highlighted match and exit the search. If zero, they will exit the search and then run, deleting a single character.

`search-in-menu` Preference Default: 0

This variable controls what Epsilon does when you press “?” during completion and then continue typing a response. If zero, Epsilon moves from the pop-up list of responses back to the prompt area, and editing keys like <Left> navigate in the response. If nonzero, Epsilon moves in the pop-up menu of names to the first name that matches what you’ve typed, and stays in the pop-up window. (If it can’t find a match, Epsilon moves back to the prompt as before.)

`search-in-region` Preference Default: 0

This variable contains bit flags that alter how searching behaves. If the 0x1 bit is set, whenever a region is highlighted when you invoke a searching command, searching will be restricted to that region, as if you had used the `search-region` command. If the 0x2 bit is set, a `grep` of a single buffer with a highlighted region will restrict itself to that region.

`search-man-pages-shows-all` Preference Default: 0

By default, the `search-man-pages` command only shows the first match in each man page. If this variable is nonzero, it shows all matches in each man page.

`search-positions-at-start` Preference Default: 0

If nonzero, nonincremental forward searching goes to the start of the match it finds, not its end.

`search-wraps` Preference Default: 1

Bits in this variable control whether searching wraps to the other end of the buffer when no more matches are found.

By default, when an incremental search fails, pressing `Ctrl-s` or `Ctrl-r` to continue the search in the same direction makes Epsilon wrap to the other end of the buffer and continue searching from there. The 1 bit enables this wrapping behavior.

The 2 bit controls whether Epsilon wraps during non-incremental searching, such as in the commands `string-search` or `search-again`. If set, Epsilon prompts when there are no more matches in the current direction and offers to wrap and search again. If not set, it simply reports no more matches.

`see-delay` Preference Default: 100

Epsilon displays most messages in the echo area for at least `see-delay` hundredths of a second before replacing them with new messages.

`selectable-colors` Default: varies

This variable contains the maximum number of color combinations the `set-color` command lets you select from.

`selected-color-scheme` Default: varies

Epsilon keeps the name table index of the current color scheme in this variable.

`sentence-end` Preference Default: [Omitted]

Epsilon uses this regular expression pattern to find the end of a sentence.

`sentence-end-double-space` Preference Default: 0

Set this variable to one if you want filling commands and sentence commands to use two spaces at the ends of sentences instead of one.

<code>server-raises-window</code>	Preference	Default: 1
-----------------------------------	------------	------------

Under X11, the `-add` and `-wait` flags cause the server instance of Epsilon to try to raise itself in the window order and set the input focus to itself, as Epsilon does under MS-Windows, if this variable is nonzero. Some window managers for X will keep programs from altering the window order in this way. Under macOS, these flags normally ask macOS to activate the XQuartz app and move it to the front, as well as asking XQuartz to make Epsilon the front window. Setting this variable to 2 makes Epsilon only ask XQuartz.

<code>session-always-restore</code>	Preference	Default: 1
-------------------------------------	------------	------------

If nonzero, Epsilon reads a session file when starting even if its command line contains file names. If zero, Epsilon only restores the previous session when no files are specified. Whenever its command line contain a file name, it doesn't read or write a session file.

<code>session-default-directory</code>	Preference	Default: none
--	------------	---------------

If Epsilon finds no session file by searching the current directory tree, it uses a session file in this directory. But if `session-default-directory` is empty, Epsilon uses the `EPSPATH` configuration variable, if there is one, or the customization directory.

<code>session-file-name</code>	Preference	Default: none
--------------------------------	------------	---------------

If this variable is nonempty, it provides the name of the session file Epsilon should use. If the name isn't an absolute pathname, Epsilon can search for files by that name in the current directory hierarchy.

<code>session-restore-biggest-file</code>	Preference	Default: 20,000,000
---	------------	---------------------

To prevent excessive delays when starting, Epsilon won't automatically restore any files bigger than this size in bytes when it restores a previous session. This variable applies to local or network files only, not files accessed with a URL, see `session-restore-biggest-remote-file` for those. If this value is negative, Epsilon ignores it and restores files of any size.

<code>session-restore-biggest-remote-file</code>	Preference	Default: 0
--	------------	------------

To prevent excessive delays when starting, Epsilon won't automatically restore any files bigger than this size in bytes when it restores a previous session. This variable applies to files accessed with URL, see `session-restore-biggest-file` for files that don't use a URL. If this value is negative, Epsilon ignores it and restores files of any size.

<code>session-restore-directory</code>	Preference	Default: 2
--	------------	------------

When Epsilon reads a session file, it can restore the current directory named in that file. If `session-restore-directory` is 1, it always does this. If 0, it never restores the current directory. If 2, the default, it restores the current directory only if the `-w1` flag was specified.

`session-restore-directory-buffers` Preference Default: 1

Bits in this variable control whether Epsilon records various types of dired buffers in a session. If 0, it doesn't record any.

The value 1 makes it records only dired buffers that don't involve URLs or use the `**` syntax to search an entire hierarchy. The value 2 makes it records dired buffers that use a URL. The value 4 makes it records dired buffers that use the `**` syntax to search an entire hierarchy. (Restoring such dired buffers can make Epsilon slow to start up.)

Add these bit values together to record dired buffers in more than one category. Also see the `session-restore-max-directories` variable.

`session-restore-exclude-file-types` Preference Default: 0

When restoring a session, bits set in this variable tell Epsilon for Windows to skip restoring files that are on certain types of drives. Set the `FSYS_NETWORK` (32) bit to keep Epsilon from trying to restore any file on a drive mounted from another computer, `FSYS_REMOVABLE` (64) to skip files from a flash drive or similar, and `FSYS_CDROM` (128) to skip files from a CD-ROM or similar. (Devices report their own types to Windows, and hardware makers aren't always consistent about devices that don't fall neatly into this scheme.)

`session-restore-files` Preference Default: 1

If 0, when Epsilon restores a session, it won't load any files (or directories) named in the session, only settings like previous search strings and command history. If 1, Epsilon will restore previous files as well as other settings. If 2, Epsilon will restore previous files only if there were no files specified on Epsilon's command line.

`session-restore-max-directories` Preference Default: 4

When Epsilon records a session, it will only remember up to this number of dired buffers, starting with the most recently accessed ones. If this value is negative, Epsilon ignores it and records any number of dired buffers.

`session-restore-max-files` Preference Default: 25

When Epsilon records a session, it will only remember up to this number of files to be reloaded the next time you start Epsilon. Files are prioritized by time of access in Epsilon, so Epsilon by default restores the 25 files you've most recently edited. If this value is negative, Epsilon ignores it and restores any number of files.

`session-tree-root` Preference Default: "NONE"

If nonempty, when Epsilon searches for a session file in the current directory tree, it only examines directories that are children of this directory. For example, if `session-tree-root` holds `\joe\proj`, and the current directory is `\joe\proj\src`, Epsilon will search in `\joe\proj\src`, then `\joe\proj`, for a session file. If the current directory is `\joe\misc`, on the other hand, Epsilon won't search at all (since `\joe\misc` isn't a child of `\joe\proj`), but will use the rules in the previous paragraph. By default, `session-tree-root` is set to an impossible absolute pathname, so searching is disabled.

`session-warn-when-saving` Preference Default: 1

On exit, Epsilon tries to save the current session. If an error occurs when saving the session, it asks for confirmation before exiting. Set this variable to zero to disable that prompt.

`set-bookmark-automatically` Preference Default: 0

If this variable is 1, after every search command, Epsilon will set a bookmark with your previous buffer location, like running the `set-bookmark` command.

`shell-auto-show-delim-chars` Preference Default: "[()]"

This variable holds the set of delimiter characters that should trigger Epsilon's auto-show-delimiters feature in Shell mode. Epsilon will search for and highlight the match of each delimiter.

`shell-heredoc-syntax` Preference Default: 0

This variable controls how Epsilon syntax highlights a here document in Shell mode. If the heredoc's terminating string ends in ".html" or ".tex", Epsilon always colors it using the corresponding mode. Otherwise, it checks this variable. The bit 0x100 makes it try to autodetect TeX code, and the bit 0x200 makes it try to autodetect HTML code. If neither is detected, the bit 1 makes it default to TeX, and 2 makes it default to HTML. (Only include one of these.) If neither bit is present, the here document is colored with the perl-constant color class. (Shell mode uses the same color classes as Perl mode.)

`shell-tab-override` Preference Default: 8

If you want the width of a tab character in Shell script buffers to be different than in other buffers, set this variable to the desired value. Shell mode will change the buffer's tab size to the specified number of columns. Setting this variable doesn't change existing buffers; set the `tab-size` variable for that.

`shift-selecting` System Default: 0

Epsilon uses this variable to keep track of whether the currently highlighted selection was begun by pressing an arrow key while holding down the Shift key. The value 1 represents this state. In it, pressing an arrow key without holding down the Shift key will turn off highlighting. The value 2 is used for certain types of regions whose highlighting should be turned off in response to an unshifted movement key. A value of 0 means movement keys won't do anything special regarding highlighting.

`shift-selects` Preference Default: 1

If this variable is nonzero, you can select text by using the arrow keys, `<Home>`, `<End>`, `<PageUp>`, or `<PageDown>` while holding down the Shift key.

`show-all-variables` System Default: 0

If zero, commands that offer completion on variable names will only recognize user variables, those marked with the `user` keyword. If nonzero, such commands also list system variables.

`show-mouse-choices` System Default: 0

If nonzero, commands that provide completion immediately display a list of possible choices, when run via the mouse.

`show-spaces` Preference Buffer-specific Default: 0

Bits in this variable control whether Epsilon displays special symbols on the screen for each `<Space>`, `<Tab>`, or `<Newline>` character in the buffer, to make them easily visible. The value 1 enables this use of special symbols.

Other bits may be used to disable this feature for individual characters. The values 2, 4, and 8 keep Epsilon from using its special symbols for `<Space>`, `<Tab>`, and `<Newline>`, respectively. Epsilon uses a special symbol only if the 1 bit is present and that particular character's disabling bit is not.

`show-status` Default: 1

If nonzero, Epsilon displays progress messages during certain lengthy operations like sorting. Otherwise, no status messages appear.

`show-tag-line` Preference Default: 2

When Epsilon jumps to a tag, it positions the window so the first line of the definition appears this many lines from the top of the window.

`show-when-idle` Preference Default: none

You can set Epsilon to display text in the echo area whenever it's idle. The `show-when-idle` variable holds the text to display. It can include any of the following sequences, and Epsilon will substitute the indicated value for that sequence:

%c Epsilon substitutes the current column number, counting columns from 0.

%C Epsilon substitutes the current column number, counting columns from 1.

%d Epsilon substitutes the current display column, with a `<` before it, and a space after. However, if the display column has a value of 0 (meaning horizontal scrolling is enabled, but the window has not been scrolled), or -1 (meaning the window wraps long lines), Epsilon substitutes nothing.

%D Epsilon substitutes the current display column, but if the display column is -1, Epsilon substitutes nothing.

%f Epsilon substitutes the name of the current function, class, or similar (in buffers where Epsilon can determine this).

%l Epsilon substitutes the current line number. (Also see the `draw-line-numbers` variable to display line numbers to the left of text.)

%m Epsilon substitutes the text " More ", but only if characters exist past the end of the window. If the last character in the buffer appears in the window, Epsilon substitutes nothing.

%P Epsilon substitutes the percentage of point through the buffer, followed by a percent sign.

%p Epsilon substitutes the percentage of point through the buffer, followed by a percent sign. However, if the bottom of the buffer appears in the window, Epsilon displays Bot instead (or End if point is at the very end of the buffer). Epsilon displays Top if the top of the buffer appears, and All if the entire buffer is visible.

%s Epsilon substitutes “* ” if the buffer’s modified flag has a nonzero value, otherwise nothing.

%S Epsilon substitutes “*” if the buffer’s modified flag has a nonzero value, otherwise nothing.

%h Epsilon substitutes the current hour in the range 1 to 12.

%H Epsilon substitutes the current hour in military time in the range 0 to 23.

%n Epsilon substitutes the current minute in the range 0 to 59.

%e Epsilon substitutes the current second in the range 0 to 59.

%a Epsilon substitutes “am” or “pm” as appropriate.

Note: For the current time, use a sequence like `%2h:%02n %a` for “3:45 pm” or `%02H:%02n:%02e` for “15:45:21”.

% % Epsilon substitutes a literal “%” character.

For any numeric substitution, you may include a number between the % and the letter code, giving the field width: the minimum number of characters to print. You can use the same kinds of field width specifiers as C’s `printf()` function. The sequence `%4c` expands to “<Space><Space><Space>9”, `%04c` expands to “0009”, and `%-4c` expands to “9<Space><Space><Space>”.

Also see the `mode-format` variable.

<code>show-when-idle-column</code>	Preference	Default: 48
------------------------------------	------------	-------------

You can set Epsilon to display text in the echo area whenever it’s idle. Epsilon positions the text `show-when-idle-column` columns from the left edge of the screen. Set this variable to a negative number to make Epsilon count columns from the right edge of the screen instead. For example, set `show-when-idle-column` to -10 to make Epsilon position the text 10 columns from the right edge.

<code>soft-tab-size</code>	Preference Buffer-specific	Default: 0
----------------------------	-------------------------------	------------

If nonzero, indenting commands like `indent-rigidly` and `back-to-tab-stop` will indent by this amount instead of the setting of the `tab-size` variable.

If this variable is zero in a buffer, and there’s a variable named `thismode-soft-tab-size`, where *thismode* is the current mode’s name, Epsilon uses that instead of the default value of `soft-tab-size`.

<code>sort-case-fold</code>	Preference Buffer-specific	Default: 2
-----------------------------	-------------------------------	------------

When comparing lines of text during sorting, Epsilon folds lower case letters to upper case before comparison, if the `sort-case-fold` variable is 1. If the `sort-case-fold` variable is 0, Epsilon compares characters as-is. If `sort-case-fold` is 2, Epsilon instead folds characters only if the `case-fold` variable is nonzero.

`spell-helper-program`

Preference

Default: ""

The `spell-correct` command uses this variable when it generates a list of suggestions. If it's empty, Epsilon uses its own internal logic for guessing the correct word. If it's nonempty, it must be a command line, which Epsilon will execute. The program that runs should engage in a dialog following the “`ispell -a`” model, in which Epsilon sends it words and it replies with suggestions. Use the `spell-configure` command to set this variable.

Any `%x` in the variable will be replaced with the name of the directory containing Epsilon's executable; any `%X` will be replaced by the 8.3 version of that directory name.

`ssh-command-windows`

Preference

Default: `%w ssh`

When Epsilon evaluates the `ssh-template` and `ssh-no-user-template` variables, it substitutes “`ssh`” for any `%c` sequence it finds. But an OS-specific variable like this one can override this. Under Windows, Epsilon uses this variable as the name of the `ssh` command.

Any `%w` in the setting will be replaced by an invocation of an appropriate helper program when running on a system with Cygwin installed (or with nothing if Cygwin is not installed). (The helper program instructs any programs that follow on the command line to behave as if running interactively, when sending output or prompting for user input. In particular, this makes it prompt for a password or passphrase.)

Any `%x` in the setting will be replaced with the name of the directory containing Epsilon's executable; any `%X` will be replaced by the 8.3 version of that directory name.

`ssh-interpret-output`

Preference

Default: `0xfef`

In an `ssh` buffer, Epsilon can look for certain escape sequences and cursor positioning, and interpret them, translating them into coloring for instance. Bits in this variable say which sorts of sequences to look for; add them to select the rules you want.

The value `0x1` makes Epsilon look for underlining that uses backspacing, such as `_ Ctrl-h k` to produce an underlined `K`; this is common in man pages and similar.

The value `0x2` makes Epsilon look for bare `Ctrl-m` characters used to indicate overtyping and remove the text to be overwritten. Some progress messages use this technique.

The value `0x4` makes Epsilon look for ANSI escape sequences that generate colors. Epsilon only recognizes certain patterns of escape sequences that color specific sections of output.

The value `0x8` makes Epsilon look for some uses of backspacing to produce bold characters.

If you've scrolled back in the buffer, Epsilon can jump to the end of the buffer whenever new output from the remote system arrives to show you the new output. Set the `0x10` bit to enable this.

The value `0x20` makes Epsilon delete any NUL characters or characters with ASCII code 1 from `ssh` output.

The value `0x40` makes Epsilon delete certain sequences defined for the `xterm` terminal emulator.

The value `0x80` makes Epsilon look for certain uses of backspacing to overwrite text with new text.

The value `0x100` makes Epsilon interpret Unicode UTF-8 encoding in the output.

<code>ssh-no-user-template</code>	Preference	Default: "%c %h"
-----------------------------------	------------	------------------

When the `ssh` command tries to run an external `ssh` program to connect securely to another computer, it uses this template to build the command line to run. Epsilon substitutes the host name for `%h`. For `%c`, Epsilon substitutes the contents of the variable `ssh-command-windows` under Windows, and “`ssh`” elsewhere. Any `%x` in the setting will be replaced with the name of the directory containing Epsilon’s executable; any `%X` will be replaced by the 8.3 version of that directory name.

If the `ssh` command contains a user name, Epsilon uses the `ssh-template` variable instead. Epsilon checks for host-specific variables before checking this variable.

<code>ssh-template</code>	Preference	Default: "%c -l %u %h"
---------------------------	------------	------------------------

When the `ssh` command tries to run an external `ssh` program to connect securely to another computer, it uses this template to build the command line to run. Epsilon substitutes the host name for `%h` and the user name for `%u`. For `%c`, Epsilon substitutes the contents of the variable `ssh-command-windows` under Windows, and “`ssh`” elsewhere. Any `%x` in the setting will be replaced with the name of the directory containing Epsilon’s executable; any `%X` will be replaced by the 8.3 version of that directory name.

If the `ssh` command contains no user name, Epsilon uses the `ssh-no-user-template` variable instead. Epsilon checks for host-specific variables before checking this variable.

<code>start-make-in-buffer-directory</code>	Preference	Default: 2
---	------------	------------

The `start-make-in-buffer-directory` variable controls which directory becomes current when you run the `make` command. Set the variable to 0 if you want each subprocess to begin with its current directory set to match Epsilon’s. Set the variable to 2 if you want each subprocess to begin in the current buffer’s directory. Set the variable to 1 if you want each subprocess to begin in the current buffer’s directory, and you also want Epsilon to change its own current directory to match, whenever you start a process. Also see the `start-process-in-buffer-directory` variable.

<code>start-process-in-buffer-directory</code>	Preference	Default: 2
--	------------	------------

The `start-process-in-buffer-directory` variable controls which directory becomes current when you start a process. Set the variable to 0 if you want each subprocess to begin with its current directory set to match Epsilon’s. Set the variable to 2 if you want each subprocess to begin in the current buffer’s directory. Set the variable to 1 if you want each subprocess to begin in the current buffer’s directory, and you also want Epsilon to change its own current directory to match, whenever you start a process. Also see the `start-make-in-buffer-directory` variable.

<code>state-extension</code>	System	Default: ".sta"
------------------------------	--------	-----------------

This variable holds the correct extension of state files in this version of Epsilon.

<code>state-file-backup-name</code>	Preference	Default: "%pebackup%e"
-------------------------------------	------------	------------------------

When you write a new state file, Epsilon makes a copy of the old one if the variable `want-state-file-backups` is nonzero. Epsilon constructs the backup file name from the original using the file name template in this variable.

<code>switch-buffers-options</code>	Preference	Default: 0
-------------------------------------	------------	------------

The `switch-buffers` command on Ctrl-Tab normally displays buffers by order of last access, with the current buffer at the top. Bits in this variable change that.

The 1 bit makes `switch-buffers` use the ordering set for the `bufed` command. Change it by pressing A, B, F, I, or U in a `bufed` listing.

Set the 2 bit to make `switch-buffers` use the `bufed-grouping` variable to control buffer order, listing modified buffers separately from unmodified buffers, and buffers with files separately from those without one. Without this bit, `switch-buffers` sorts all these types of buffers together, regardless of `bufed-grouping`.

<code>system-window</code>	System	Window-specific	Default: 0
----------------------------	--------	-----------------	------------

If nonzero in a window, user commands that switch windows will skip over this window.

<code>tab-size</code>	Preference	Buffer-specific	Default: 8
-----------------------	------------	-----------------	------------

This variable holds the number of columns from one tab stop to the next. Epsilon expands tab characters in the buffer to reach the next tab stop. By default, Epsilon also indents in units of the tab size. Set the `soft-tab-size` variable if you want independent settings for the width of a tab character and the amount to indent.

<code>table-count</code>	System	Default: 0
--------------------------	--------	------------

This variable counts the number of prefix keys like Ctrl-x you've typed so far in the current command.

<code>tag-ask-before-retagging</code>	Preference	Default: 0
---------------------------------------	------------	------------

If zero, when a tag's line has changed within a file, Epsilon retags the file automatically and then searches again. Similarly, when Epsilon can't find a tag at all, it tries tagging the current file. If nonzero, Epsilon asks before doing either of these things.

<code>tag-batch-mode</code>	System	Default: 0
-----------------------------	--------	------------

Epsilon's tag facility uses this variable to decide if it should report an error immediately, or just log it to a buffer.

<code>tag-by-text</code>	Preference	Default: 1
--------------------------	------------	------------

If nonzero, Epsilon includes the entire line that defined a tag in the tag file, so it can search for the line when the buffer has been modified since tagging. If zero, Epsilon only includes the offset, saving space in the tag file for files that rarely change.

<code>tag-c-preprocessor-skip-pat</code>	Preference	Default: "0"
--	------------	--------------

When tagging in C mode, Epsilon skips over `#if 0` blocks. It can skip other conditionals too. This variable should contain a regular expression, like `(0|DEBUG)`; if what follows `#if` matches it, that block will be skipped.

tag-case-sensitive	Preference	Default: 0
--------------------	------------	------------

Set this variable nonzero if you want tagging to consider MAIN, Main and main to be distinct tags. By default, typing main will find any of these.

tag-declarations	Preference	Default: 0
------------------	------------	------------

The tag-declarations variable lets you set whether the tagger will tag function or variable declarations (as opposed to definitions, which Epsilon always tags). If zero (the default), Epsilon only tags definitions. If one, Epsilon tags function declarations as well. If two, Epsilon tags variable declarations (which use the `extern` keyword). If three, Epsilon tags both types of declarations. You may wish to use this setting to tag the `.h` header files of library functions.

tag-display-width	Preference	Default: 40
-------------------	------------	-------------

When Epsilon lists tags during tag completion, it shows the tag name followed by its file name. This variable controls the width available for the tag name. It should be bigger than the longest expected tag name.

tag-list-exact-only	System	Default: 0
---------------------	--------	------------

Epsilon's tag facility uses this variable internally to decide if tag matching should include prefix matches or only exact matches.

tag-modifier-kwd	System	Default: 0
------------------	--------	------------

The C tagger uses this variable internally to keep track of whether it's found keywords like `extern`, `typedef`, or `import` in the current declaration.

tag-options	Preference	Default: 2
-------------	------------	------------

This variable contains bits that modify how tagging commands behave. The 1 bit makes the `tag-files` command delete all current tags before finding new tags. The 2 bit makes the `select-tag-file` command look up the directory hierarchy from the current file when suggesting a new tag file.

tag-pattern-c	System	Default: [Omitted]
---------------	--------	--------------------

The `pluck-tag` command searches using this regular expression to locate the current tag in C/C++/Java buffers.

tag-pattern-default	System	Default: [a-zA-Z0-9_]+
---------------------	--------	------------------------

The `pluck-tag` command searches using this regular expression to locate the current tag in buffers without a mode-specific tag pattern.

tag-pattern-perl	System	Default: [Omitted]
------------------	--------	--------------------

The `pluck-tag` command searches using this regular expression to locate the current tag in Perl buffers.

<code>tag-relative</code>	Preference	Default: 1
If nonzero, Epsilon stores relative pathnames in the tag file whenever it can. If zero, Epsilon uses only absolute pathnames.		
<code>tag-show-percent</code>	System	Default: 0
If nonzero, Epsilon displays a percentage status report while tagging instead of mentioning each tag it finds. Commands that use tagging to parse a buffer without really generating tags can set this.		
<code>tag-which-items</code>	Preference	Default: 0xff
Bits in this variable determine which types of items are tagged. The value 0x1 includes functions, 0x2 includes variables, 0x4 includes macros, 0x8 includes the values of enums, 0x10 includes the names of classes, structures, unions, and similar items, and 0x20 includes the names of enums. Also see <code>list-which-definitions</code> . Some modes may not distinguish all these types of items.		
<code>tcl-auto-show-delim-chars</code>	Preference	Default: "[()]"
This variable holds the set of delimiter characters that should trigger Epsilon's auto-show-delimiters feature in Tcl mode. Epsilon will search for and highlight the match of each delimiter.		
<code>tcl-indent</code>	Preference Buffer-specific	Default: 4
Tcl mode indents by this many columns for each additional level of nesting.		
<code>telnet-interpret-output</code>	Preference	Default: 0xfef
In a Telnet buffer, Epsilon can look for certain escape sequences and cursor positioning, and interpret them, translating them into coloring for instance. Bits in this variable say which sorts of sequences to look for; add them to select the rules you want.		
The value 0x1 makes Epsilon look for underlining that uses backspacing, such as <code>_ Ctrl-h k</code> to produce an underlined K; this is common in man pages and similar.		
The value 0x2 makes Epsilon look for bare Ctrl-m characters used to indicate overtyping and remove the text to be overwritten. Some progress messages use this technique.		
The value 0x4 makes Epsilon look for ANSI escape sequences that generate colors. Epsilon only recognizes certain patterns of escape sequences that color specific sections of output.		
The value 0x8 makes Epsilon look for some uses of backspacing to produce bold characters.		
If you've scrolled back in the buffer, Epsilon can jump to the end of the buffer whenever new output from the remote system arrives to show you the new output. Set the 0x10 bit to enable this.		
The value 0x20 makes Epsilon delete any NUL characters or characters with ASCII code 1 from telnet output.		
The value 0x40 makes Epsilon delete certain sequences defined for the xterm terminal emulator.		
The value 0x80 makes Epsilon look for certain uses of backspacing to overwrite text with new text.		
The value 0x100 makes Epsilon interpret Unicode UTF-8 encoding in the output.		

`tex-auto-fill-mode` Default: 1

If nonzero, Epsilon breaks long lines in TeX/LaTeX files using auto-fill mode. If zero, it doesn't.

`tex-auto-show-delim-chars` Default: "[{()}]"

This variable holds the set of delimiter characters that should trigger Epsilon's auto-show-delimiters feature in TeX mode. Epsilon will search for and highlight the match of each delimiter.

`tex-environment-name` Default: "document"

The `tex-environment` command uses this variable to hold the name of the last environment you inserted in TeX mode.

`tex-force-latex` Preference Buffer-specific Default: 1

Some TeX mode commands are slightly different in LaTeX than in pure TeX. Set `tex-force-latex` to 1 if all your documents are LaTeX, 0 if all your documents are TeX, or 2 if Epsilon should determine this on a document-by-document basis. In that case, Epsilon will assume a document is LaTeX if it contains a `\begin{document}` statement or if it's in a file with an `.ltx` extension.

`tex-look-back` Preference Default: 20000

TeX syntax highlighting sometimes needs to look back in the buffer to locate the start of a paragraph. Long stretches of text without paragraph breaks can make it slow. Set this variable lower if you want Epsilon to give up sooner and incorrectly color some rare cases.

`tex-paragraphs` Preference Buffer-specific Default: 0

If nonzero, then Epsilon will not consider as part of a paragraph any sequence of lines that each start with at sign or period, if that sequence appears next to a blank line. And lines starting with `\begin` or `\end`, or with `%`, `\[`, `\]`, or `$$`, or ending with `\`, will also delimit paragraphs.

`tex-save-new-environments` Preference Default: 0

The `tex-environment` command lets you easily create a new environment, inserting `begin/end` pairs. When it prompts for an environment name, you can type the name of a new environment, and Epsilon will remember it for the rest of the editing session, offering it for completion. Set this variable nonzero and Epsilon will also save the new environment name for future sessions.

`tex-spell-options` Preference Default: 0

This variable controls the Spell minor mode in TeX mode. Use the `spell-mode` command to set it to 1, and Epsilon will highlight misspelled words in TeX text. See the `default-spell-options` variable for the other bits you can set to customize spell checking in TeX mode.

`text-color` Window-specific Default: 0

This variable contains the color class of normal text in the current window.

`this-cmd` Default: none

Some commands behave differently depending on what command preceded them. To get this behavior, the command acts differently if `prev-cmd` is set to a certain value and sets `this-cmd` to that value itself. Epsilon copies the value in `this-cmd` to `prev-cmd` and then clears `this-cmd` each time through the main loop.

`tiled-border` Preference Default: 0xAA

This variable holds the border codes Epsilon uses for putting borders at the edges of a tiled window.

`tiled-scroll-bar` System Default: 0

If nonzero, Epsilon constantly displays a scroll bar on tiled windows. Set this with the `toggle-scroll-bar` command.

`topindent` Preference Buffer-specific Default: 1

If nonzero, Epsilon indents top-level statements in a function. If zero, Epsilon keeps such statements at the left margin.

`translation-type` Buffer-specific Default: 5

Epsilon uses this variable to record the type of line translation and Unicode encoding used by the current buffer. The `set-line-translate` and `set-encoding` commands set this variable. To read a new file in a mode other than the default, type `Ctrl-u Ctrl-x Ctrl-f` to run the `find-file` command with a numeric argument.

`type-point` Buffer-specific Default: none

This variable holds the position within the process buffer where Epsilon inserts new text from the process. Epsilon retrieves any text after the type point and sends it as input to the process. Other functions that run in the background and insert retrieved text in a buffer, such as Epsilon's internet features, also use this variable.

`typing-deletes-highlight` Preference Default: 1

If this variable is 1, pressing a self-inserting key like “j” while text is highlighted deletes the highlighted selection, replacing it with the key. Pressing (Backspace) simply deletes the text. If this variable is 0, a highlighted region isn't deleted. If this variable is 2, deleting highlighted regions happens only at prompts.

If you set this variable to zero, you may wish to set the `insert-default-response` variable to zero also. Then Epsilon won't automatically insert and highlight your previous response at various prompts.

`typing-hides-highlight` Preference Default: 1

If this variable is nonzero, pressing a self-inserting key like “j” while text is highlighted turns off the highlighting. Setting this variable to zero has little effect when `typing-deletes-highlight` is set to delete highlighted regions.

`uncompress-files` Preference Default: 3

If this variable is 1, reading a file with a `.gz` extension will uncompress it and display the uncompressed version. If this variable is 2, Epsilon will do the same for files with a `.bz2` extension. A value of 3 does both, and 0 does neither.

This feature runs the standard utility programs `gzip` (for `.gz`) and `bzip2` (for `.bz2`); you'll need to install them if they're not already installed. For Windows, the programs provided by the Cygwin environment will work fine.

`undo-flag` Buffer-specific Default: none

In addition to buffer changes and movements, Epsilon can record other information in its list of undoable operations. Each time you set this variable, Epsilon inserts a "flag" in its undo list with the particular value you specify. When Epsilon is undoing or redoing and encounters a flag, it immediately ends the current group of undo operations, returns a special code, and puts the value of the flag it encountered back into the `undo_flag` variable.

`undo-keeps-narrowing` System Buffer-specific Default: 0

If you use the `narrow-to-region` command to hide part of the buffer, and then use the `undo` command to undo a change in a hidden part of the buffer, `undo` removes the narrowing. Some modes set this variable `nonzero` to prevent that behavior.

`undo-size` Preference Buffer-specific Default: 500000

Epsilon retains at most this many characters of deleted or changed text in this buffer's undo information.

`ungot-key` Default: -1

If this variable is set to some value other than its usual value of -1, Epsilon uses that value when it next tries to read a key and sets `ungot-key` to -1 again.

`use-c-macro-rules` Preference Default: 2

This variable controls whether C mode indenting recognizes certain macro names used by various Microsoft development environments. Epsilon provides special indenting rules for them. The value 0 disables such special rules; 1 enables them unconditionally, and 2 enables them only in source file types that typically use such macros, and then only under Windows.

The indenting rules use the variables `c-block-macro-close-pat`, `c-block-macro-inner-pat`, and `c-block-macro-open-pat`.

`use-compile-command-file-variable` Preference Default: 2

A file can use a file variable named "compile-command" to indicate that the `compile-buffer` command should run a specific command line to compile that file. Set this variable to 1 if you want `compile-buffer` to use such settings without any further prompting. Set it to 0 to have `compile-buffer` ignore any such file variable. Set it to 2 to have the command prompt, showing you the suggested compilation command and letting you edit it, before running it.

`use-default` System Default: 0

If nonzero, every time Epsilon refers to a buffer- or window-specific variable, it uses the default value instead of the current value.

`use-file-variables` Preference Default: 7

When Epsilon reads a file, it looks for special lines that can customize certain buffer-specific settings like the tab size or fill column appropriately for that file, or specify the correct mode for the file. These settings are called file variables. Bits in the `use-file-variables` variable control this. See page 127 for details.

The 1 bit lets Epsilon scan each file it reads for file variables in its native (Emacs-compatible) format.

The 2 bit lets Epsilon first look for file variables in a special file named `.epsilon_vars` that applies to all files in a directory. Before that, this bit also tells it to load any file variables it finds in a file named `.epsilon_vars` in its customization directory.

The 4 bit lets Epsilon scan each file for file variables in the format used by the vi/vim family of editors. Vi documentation uses the name “modelines” for what Epsilon calls file variables.

The 8 bit lets Epsilon look for `.epsilon_vars` files even during `grep`. By default it doesn’t, for speed reasons.

The 16 bit lets Epsilon scan up a directory tree for an `.epsilon_vars` file, if there isn’t one in a given file’s directory. By default it doesn’t.

`use-grep-ignore-file-variables` Preference Default: 1

Set `use-grep-ignore-file-variables` to zero if you want Epsilon to ignore the various `grep-ignore-` variables, and search all files.

`use-process-current-directory` Preference Default: 1

Epsilon tries to retrieve the process’s current directory and use it as the default directory for the process buffer.

If this variable is set to 2, Epsilon for Windows will also modify its own current directory (set with `Alt-x cd`) when it detects that a process’s current directory has changed.

`user-abort` Default: 0

Epsilon sets this nonzero when you press the abort key. Commands check this variable and abort if it’s nonzero.

`version` Default: varies

This variable holds the current version number of Epsilon in text form, as recorded in the executable file.

`vbasic-auto-show-delim-chars` Preference Default: "`[]`"

This variable holds the set of delimiter characters that should trigger Epsilon’s `auto-show-delimiters` feature in Visual Basic mode. Epsilon will search for and highlight the match of each delimiter.

`vbasic-indent` Preference Buffer-specific Default: 3

Each level of indentation in Visual Basic mode will occupy this many columns.

`vbasic-indent-case` Preference Default: 1

If this variable is nonzero, Case statements in Select blocks in Visual Basic will receive extra indentation. If zero, Case statements will be indented to the same level as their Select blocks.

`vbasic-indent-subroutines` Preference Default: 1

If nonzero, the bodies of subroutines will be indented more than the subroutine declaration line at the top, in Visual Basic mode. Otherwise they will start with the same indentation.

`vbasic-indent-with-tabs` Preference Default: 0

If zero, Epsilon indents using only space characters, not tab characters, in Visual Basic mode. The `vbasic-mode` command initializes the `indent-with-tabs` variable from this one.

`vbasic-language-level` Preference Default: 0xffff

When Epsilon colors Visual Basic files, it uses bits in this variable to control which set of keywords to color.

The 0x2 bit makes it color keywords added in VB.NET 2005.

The 0x4 bit makes it color keywords added in VB.NET 2008.

The 0x40 bit makes it color unreserved VB.NET keywords like `isfalse` and `until`.

The 0x80 bit makes it color obsolete VB.NET keywords like `gosub` and `variant`.

Reload each Visual Basic file or toggle its `want-code-coloring` variable after setting this variable so the new rule takes effect.

`vbasic-reindent-previous-line` Preference Default: 1

This variable controls whether Epsilon reindents the previous line when you press `(Enter)` in Visual Basic mode.

`versioned-file-string` System Default: varies

This variable holds Epsilon's version number, formatted so that it can be part of a directory name. Epsilon for Unix looks for its configuration files in a directory whose name is built from this string; it also checks this against the variable `eel-version` to detect version mismatches between an Epsilon executable and its state file commands.

`vhdl-auto-show-delim-chars` Default: "()"

This variable holds the set of delimiter characters that should trigger Epsilon's `auto-show-delimiters` feature in VHDL mode. Epsilon will search for and highlight the match of each delimiter.

<code>vhdl-indent</code>	Preference Buffer-specific	Default: 4
--------------------------	------------------------------	------------

Each level of indentation in VHDL mode will occupy this many columns.

<code>virtual-insert-cursor</code>	Preference	Default: 93099
------------------------------------	------------	----------------

Epsilon uses the cursor shape code specified by this variable whenever the cursor is in virtual space (between characters) and Epsilon's overwrite mode is off. See the description of `normal-cursor` for details. See `virtual-insert-gui-cursor` for the Windows or X11 equivalent.

<code>virtual-insert-gui-cursor</code>	Preference	Default: 50002
--	------------	----------------

Epsilon for Windows or X11 uses the cursor shape code specified by this variable whenever the cursor is in virtual space (between characters) and Epsilon's overwrite mode is off. See the description of `normal-gui-cursor` for details.

<code>virtual-overwrite-cursor</code>	Preference	Default: 0005
---------------------------------------	------------	---------------

Epsilon uses the cursor shape code specified by this variable whenever the cursor is in virtual space (between characters) and Epsilon's overwrite mode is on. See the description of `normal-cursor` for details. See `virtual-overwrite-gui-cursor` for the Windows or X11 equivalent.

<code>virtual-overwrite-gui-cursor</code>	Preference	Default: 50100
---	------------	----------------

Epsilon for Windows or X11 uses the cursor shape code specified by this variable whenever the cursor is in virtual space (between characters) and Epsilon's overwrite mode is off. See the description of `normal-gui-cursor` for details.

<code>virtual-space</code>	Preference Buffer-specific	Default: 0
----------------------------	------------------------------	------------

If zero, Epsilon commands only position to places on the screen where there is actual buffer text. If one, the `<Up>` and `<Down>` keys will stay in the same column, even if no text exists there. If two, in addition to `<Up>` and `<Down>`, the `<Right>` and `<Left>` keys will move into places where no text exists, always remaining on the same line of the buffer.

<code>w-bottom</code>	System	Default: none
-----------------------	--------	---------------

Mouse commands store the bottom edge of the selected window here.

<code>w-left</code>	System	Default: none
---------------------	--------	---------------

Mouse commands store the left edge of the selected window here.

<code>w-right</code>	System	Default: none
----------------------	--------	---------------

Mouse commands store the right edge of the selected window here.

<code>w-top</code>	System	Default: none
--------------------	--------	---------------

Mouse commands store the top edge of the selected window here.

want-auto-save	Preference	Default: 1
<p>If nonzero, Epsilon periodically saves a copy of each unsaved file. Bits in this variable control various aspects of auto-saving. The 2 bit makes auto-saving more verbose, with a message every time Epsilon auto-saves a file. The 4 bit prevents Epsilon from deleting the auto-save file once the corresponding file has been saved. When Epsilon exits normally, if auto-saving is on, it deletes any auto-saved file associated with every file you're editing, even if it wasn't created by the current session. The 8 bit prevents this.</p>		
want-backups	Preference Buffer-specific	Default: 0
<p>If 2, Epsilon makes a backup whenever it saves a file. If 1, Epsilon makes a backup the first time it saves a file in a session.</p>		
want-bell	Preference	Default: 1
<p>If nonzero, Epsilon beeps to warn you of certain conditions. Variables starting with bell-on-permit finer control over just when Epsilon beeps.</p>		
want-code-coloring	Preference Buffer-specific	Default: 1
<p>If this buffer-specific variable is non-zero, Epsilon tries to do code coloring (syntax highlighting) in the current buffer.</p>		
want-cols	System	Default: varies
<p>This variable holds the value the user specified through the -vc switch, or 0 if the user did not explicitly specify the number of columns to display via this flag.</p>		
want-common-file-dialog	Preference	Default: 1
<p>In Epsilon for Windows, some commands that prompt for files can use the Windows Common File Dialog. By default, these commands use the dialog if you invoke them from the menu or tool bar, but not if you invoke them from the keyboard using their bindings. Set this variable to 2 if you want Epsilon to use the Common File Dialog whenever it can. Set the variable to 0 to prevent Epsilon from ever using this dialog. The default value of 1 produces the behavior described above.</p>		
want-debug-coloring	Preference Buffer-specific	Default: 0
<p>Bits in this variable control a facility to help debug code-coloring subroutines written in EEL. See page 491.</p>		
want-display-host-name	Preference	Default: 1
<p>Set this variable to 0 to keep Epsilon for Unix from displaying the computer's network node host name in the window title. Set this variable to 2 to have Epsilon use the computer's fully qualified domain name instead of the configured host name.</p>		
want-gui-menu	System	Default: 1
<p>Epsilon for Windows sets this variable to indicate whether it should display a menu bar.</p>		

want-gui-printing	Preference	Default: 1
-------------------	------------	------------

If this variable is zero, printing commands in Epsilon for Windows won't use standard Windows printing features, but instead will print via the `print-destination` variable. If you want Epsilon to run an external command to print a file, set this variable to zero.

want-gui-prompts	Preference	Default: 1
------------------	------------	------------

Bits in this variable control how Epsilon for Windows handles dialogs. The 1 bit tells Epsilon to display Windows-style dialogs for its prompts. Without it, Epsilon for Windows will avoid using Windows dialogs in many commands, and will draw text boxes instead, similar to the non-Windows versions of Epsilon.

Normally in a Windows dialog under Epsilon, the `<Tab>` key performs completion, just as in non-Windows versions of Epsilon. You can set the 2 bit to make it cycle through the dialog's controls instead, as in other Windows programs.

Similarly, in a Windows dialog under Epsilon, the `<Esc>` key normally performs completion. The 4 bit makes it cancel the current dialog instead, as in other Windows programs. (Regardless of this setting, pressing the abort key, normally `Ctrl-g`, always cancels the dialog.)

want-lines	System	Default: varies
------------	--------	-----------------

This variable holds the value the user specified through the `-vl` switch, or 0 if the user did not explicitly specify the number of lines to display via this flag.

want-sorted-tags	Preference	Default: 1
------------------	------------	------------

If nonzero, Epsilon displays its list of tags alphabetically. If zero, the order depends on the order in which you tagged the files.

want-state-file-backups	Preference	Default: 1
-------------------------	------------	------------

If nonzero, Epsilon makes a backup whenever you write a new state file.

want-toolbar	Preference	Default: 1
--------------	------------	------------

Epsilon uses this variable to remember if the user wants a tool bar displayed, in versions of Epsilon which support this. Use the `toggle-toolbar` command to change this setting.

want-warn	Preference	Buffer-specific	Default: 1
-----------	------------	-----------------	------------

If nonzero, before Epsilon saves a file, it checks the time and date of the copy of the file already on disk (to see if anyone has modified it since you read it into Epsilon), and warns you if the file has been modified. Epsilon also checks the file each time you switch to a buffer or window displaying that file, and before you read or write the file.

want-window-borders	Preference	Default: 1
---------------------	------------	------------

The `toggle-borders` command uses this variable to record whether or not you want borders between tiled windows. Without borders, Epsilon assigns separate color schemes to each window.

warn-before-overwrite Preference Default: 1

Commands like `write-region` that write to a user-specified file ask for confirmation if the file already exists. To make Epsilon write over such files without asking, set this variable to 0. If this variable is 2, Epsilon will warn for local files, but skip this check for remote files (accessed via an scp or ftp URL).

was-quoted System Default: 0

Epsilon makes this variable nonzero if the last file name you typed included the " character. Epsilon treats some files patterns differently in this case.

web-search-url Preference Default: [Omitted]

The `web-search` command appends the desired search text to this variable's contents, then has your web browser display the resulting URL.

wheel-click-lines Preference Default: -1

Rolling the wheel on a mouse scrolls by this many lines at once. A value of 0 means scroll by pages. Under Windows, a value of -1 means use the value set in the Mouse control panel; for Unix, it's the same as 3.

window-bufnum System Window-specific Default: none

This variable holds the buffer number of the buffer Epsilon should display in the current window.

window-caption Preference Default: "Epsilon"

Epsilon for Windows or X11 sets its caption to this text when the current buffer is not associated with a file and the `window-caption-buffer` variable is empty.

window-caption-buffer Preference Default: "(%s) - Epsilon"

Epsilon for Windows or X11 sets its caption to this text when the current buffer is not associated with a file. The %s in the text is replaced by the buffer name. If this variable is empty, Epsilon uses the `window-caption` variable instead.

window-caption-file Preference Default: "%s - Epsilon"

Epsilon for Windows or X11 sets its caption to this text when the current buffer is associated with a file. The %s in the text is replaced by the file name.

window-color-scheme System Window-specific Default: 0

If the window-specific variable `window_color_scheme` is non-zero in a window, Epsilon uses its value in place of the `selected_color_scheme` variable when displaying that window. But the similar `buffer_color_scheme` variable takes precedence over this one.

<code>window-end</code>	Window-specific	Default: none
On each screen refresh, Epsilon sets this variable to the last buffer position displayed in the window.		
<code>window-handle</code>		Default: none
This variable holds the current window's window handle, a code that uniquely identifies the window. Setting it switches windows.		
<code>window-height</code>	Window-specific	Default: none
This variable contains the height of the current window in lines, including any mode line or borders. Setting it changes the size of the window.		
<code>window-left</code>	Window-specific	Default: none
This variable holds the screen coordinate of the left edge of the current window. If the current window is a pop-window, you can set this variable to move the window around.		
<code>window-number</code>		Default: none
This variable holds a number that denotes the current window's position in the window order. Tiled windows are numbered from the upper-left window, which is numbered zero, to the lower-right window. Pop-up windows always come after tiled windows in this order, with the most recently created pop-up window last.		
<code>window-overlap</code>	Preference	Default: 2
When scrolling by pages, Epsilon leaves this many lines of overlap between one window of text and the next (or previous). A negative value for <code>window-overlap</code> represents a percentage of overlap, instead of the number of screen lines.		
<code>window-start</code>	Window-specific	Default: none
This variable holds the buffer position of the first character displayed in the current window.		
<code>window-top</code>	Window-specific	Default: none
This variable holds the screen coordinate of the top edge of the current window. If the current window is a pop-window, you can set this variable to move the window around.		
<code>window-width</code>	Window-specific	Default: none
This variable contains the width of the current window in characters, including any borders. Setting it changes the size of the window.		
<code>winhelp-display-contents</code>	Preference	Default: 0
If <code>winhelp-display-contents</code> is nonzero, menu items that look up text in a .hlp file will display the contents page of their help file if you select one without first highlighting a keyword. If zero, Epsilon will display the keyword index of the help file.		

<code>word-pattern</code>	Buffer-specific	Default: <code>points to default_word</code>
---------------------------	-----------------	--

This variable points to the regular-expression pattern Epsilon uses to move forward or backward by a word in the current buffer. Set the variable `default-word` instead to change the pattern for all buffers, or to change it permanently.

<code>wrap-dired-live-link</code>	Preference	Default: 1
-----------------------------------	------------	------------

When you press lowercase L in a dired buffer to create a live link window, Epsilon normally sets its window to wrap long lines. This is useful because displaying files with extremely long lines (such as binary files) can be slow when the window is set to horizontally scroll, producing a delay before Epsilon responds to keystrokes.

But you can set this variable to 0 to have Epsilon horizontally scroll long lines in such windows instead of wrapping them. Set it to 2 to maintain the current window's setting for wrapping versus horizontally scrolling, or 1 to force wrapping (the default).

<code>wrap-grep</code>	Preference	Default: 3
------------------------	------------	------------

When you use the `grep` command to list matches in files, Epsilon normally sets the current window to horizontally scroll long lines, so that long lines in the `grep` listing each occupy a single screen line. Set this variable to 1 to have it set the current window to wrap long lines instead of scrolling them. Set it to 2 to maintain the current window's setting for wrapping versus horizontally scrolling, or 0 always to force scrolling.

The default setting 3 makes Epsilon force horizontal scrolling except when the default behavior for windows has been set to wrap, in which case it maintains the current window's setting.

<code>wrap-info-mode</code>	Preference	Default: 1
-----------------------------	------------	------------

When Epsilon displays documentation in Info format, it normally sets the current window to wrap long lines. Set this variable to 0 to have it horizontally scroll long lines instead of wrapping them. Set it to 2 to maintain the current window's setting for wrapping versus horizontally scrolling, or 1 to force wrapping (the default).

<code>wrap-split-vertically</code>	Preference	Default: 0
------------------------------------	------------	------------

When you use the `split-window-vertically` command to create side-by-side windows, Epsilon normally sets them both to horizontally scroll long lines. Set this variable to 1 to have it wrap long lines instead of scrolling them. Set it to 2 to maintain the current window's setting for wrapping versus horizontally scrolling, or 0 to force scrolling (the default).

<code>wrap-view-error-lines</code>	Preference	Default: 2
------------------------------------	------------	------------

When you set the `process-view-error-lines` variable so that compiling commands display compiler error messages in a separate window, you can tell Epsilon to force that window to horizontally scroll its long lines, or to wrap them, by setting this variable to 0 or to 1, respectively. The default value 2 makes the new window inherit this setting from the adjacent window when it's created.

xml-asp-coloring Preference Default: 10

This variable tells Epsilon how to syntax highlight scripts embedded in `<% %>` delimiters in XML documents, when the file doesn't name any specific script language. Zero means use a single color, 1 means color as Javascript, 2 means color as VBScript, 3 means color as PHP, 4 means Python, 5 means CSS, and 10 means `<%` doesn't start an embedded script.

xml-auto-fill-combine Preference Default: 10

When auto-filling breaks a line of XML, it avoids breaking a line that contains just the start of an element, except when the element name is longer than specified by this variable. For example, it won't split a line right after `<type` at its start, but it may after `<newdef : address`.

xml-auto-fill-mode Preference Default: 25

This variable controls whether Epsilon automatically breaks long lines as you type in XML mode. The 1 bit toggles filling on and off entirely. It's set by the `auto-fill-mode` command. Other bits control where filling occurs.

The 2 bit lets it break text that isn't part of an XML tag. The 4 bit lets it break XML tags. The 8 bit lets it break XML comments. The value 16 lets it break comments in any embedded scripting. By default, Epsilon breaks only in comments (both types).

xml-auto-indent Preference Default: 0xff

This variable controls automatic indentation when you press `<Enter>` in XML mode. Bits in the variable control whether Epsilon auto-indents in specific regions of the document. The 0x1 bit makes Epsilon auto-indent when you press `<Enter>` outside script blocks. Other bits, as shown in the table below, make Epsilon auto-indent in that type of scripting.

Scripting Language	Bit
JavaScript	0x2
VBScript	0x4
PHP	0x8
Python	0x10
CSS	0x20

When you disable smart auto-indenting in a certain type of scripting by setting one of these bits, `<Enter>` will instead indent to the previous line. Set the 0x4000 bit if you prefer no indenting at all for regions where you've disabled smart indenting.

xml-indent Preference Buffer-specific Default: 3

Each level of indentation in XML mode will occupy this many columns. If this variable is zero, Epsilon uses the tab size instead.

xml-paragraph-elements Preference Default: "<alphanum|. | _ | hyphen>+"

When filling paragraphs in XML mode, Epsilon treats tags matching element names on this list as paragraph delimiters. The list uses the syntax of a regular expression pattern, so by default it matches all tags and only fills text between tags.

xml-php-coloring Preference Default: 10

This variable tells Epsilon how to syntax highlight scripts embedded in `<? ?>` delimiters in XML documents, when the file doesn't name any specific script language. Zero means use a single color, 1 means color as Javascript, 2 means color as VBScript, 3 means color as PHP, 4 means Python, 5 means CSS, and 10 means `<%` doesn't start an embedded script.

xml-reindent-previous-line Preference Default: 1

This variable controls whether Epsilon reindents the previous line when you press `(Enter)` in XML mode.

xml-spell-options Preference Default: 0

This variable controls the Spell minor mode in XML mode. Use the `spell-mode` command to set it to 1, and Epsilon will highlight misspelled words in XML text, ignoring element names and attributes. See the `default-spell-options` variable for the other bits you can set to customize spell checking in XML mode.

yank-line-retains-position Preference Default: 0

This variable determines how Epsilon positions point after you yank a line region when point is in the middle of a line. If nonzero, point remains where it is in the middle of the previous line. If zero, point moves to after the yanked line.

yank-options Preference Default: 0

Bits in this variable can be used to customize aspects of the `yank` command. The 1 bit alters what happens when you yank a path name starting with a path separator character `/` or `\` at a file name prompt, after a directory name itself ending in `/` or `\`. By default, Epsilon deletes the existing directory name, treating the yanked name as an absolute path. Set the 1 bit if you want Epsilon to delete merely the path separator at the end of the existing directory name, treating the yanked path name as a relative name.

yank-rectangle-to-corner Preference Default: 1

This variable determines how Epsilon positions point and mark after you yank a rectangular region. If 1, it puts point at the bottom right corner of the region, and mark at the upper left. If 2, it puts point at the upper left and mark at the lower right. If 3, it puts mark at the upper left corner, and positions point one line below the bottom left corner (Brief-style). Note that with this last style, the `yank-pop` command will not function after yanking a rectangular region.

Chapter 7

Changing Epsilon



Epsilon provides several ways for you to change its behavior. Some commands enable you to make simple changes. For example, `set-fill-column` can change the width of filled lines of text. Commands like `bind-to-key` and `create-prefix-command` can move commands around on the keyboard, and using keyboard macros, you can build simple new commands. The remaining chapters of the manual describe how to use the Epsilon Extension Language, EEL, to make more sophisticated commands and to modify existing commands.

Unless you save them, all these types of changes go away when you exit, and you must reload them the next time you run Epsilon. (But see the `record-customizations` variable.)

There are many ways to save such changes. The easiest way to save them is with the `list-customizations` command. That's the best method, unless you have a very large number of customizations. It works by adding lines to your `init.ecm` customization file, which Epsilon reads every time it starts. Or you can customize Epsilon by manually editing that file. See page 164 for details on both techniques.

But you can also save changes by storing them in a state file. This method lets Epsilon start up faster when you have hundreds of customizations, or have made many customizations using the EEL extension language. This chapter describes various ways to customize Epsilon using a modified state file.

When it starts, Epsilon reads a state file named `epsilon-v14.sta` containing all of Epsilon's initial commands, variables, and bindings. (You can use Epsilon's `-s` flag to make Epsilon load its state from some other file. For example, "`epsilon -sfilename`" loads its commands from the file `filename.sta`. The default name includes Epsilon's major version number.)

You can change Epsilon's set of commands and settings by generating a new state file with the Epsilon command `write-state` on `Ctrl-F3`. So one way to customize Epsilon is to make your changes (bind some keys, set a variable, define some macros) and use the `write-state` command to put the changes in `epsilon-v14.sta`. Your customizations will take effect each time you run Epsilon. See page 162 for more on `write-state`.

Instead of manually setting variables, then saving them in a (binary) state file, you may want to preserve your changes in a human-readable format. Commands like `list-customizations` provide one way to do that, preserving variable settings and the like in Epsilon's command file format. You can also do this using EEL.

In that case, you may find it handy to have a file that loads your changes into a fresh Epsilon, then writes the new state file automatically. The following simple EEL file, which we'll call `changes.e`, uses features described in later chapters to do just that:

```
#include "eel.h" /* Load standard definitions. */

when_loading() /* Execute this file when loaded. */
{
    want_bell = 0; /* Turn off the bell. */
    kill_buffers = 6; /* Use only 6 kill buffers. */
    load_commands("mycmds"); /* Load my new commands. */
    do_save_state("epsilon"); /* Save these changes. */
}
```

Each time you get an update of Epsilon, you can compile this program (type `ee1 changes` outside of Epsilon) and start Epsilon with its new state file (type `epsilon`). Then when you load this file (type `F3 changes` `<Enter>` to Epsilon), Epsilon will make all your changes in the updated version and automatically save them for next time.

You can change most variables as in the example above. Some variables, however, have a separate value for each buffer. Consider, for example, the tab size (which corresponds to the value of the `tab-size` variable). This variable's value can potentially change from buffer to buffer. We call this a buffer-specific variable. Buffer-specific variables have one value for each buffer plus a special value called the default

value. The default value specifies the value for the variable in a newly created buffer. A state file stores only the default value of a buffer-specific variable.

Thus, to change the tab size permanently, you must change `tab_size`'s default value. You can use the `set-variable` command to make the change, or an EEL program. The following version of `changes.e` sets the default tab size to 5.

```
#include "eel.h" /* load standard definitions */

when_loading() /* execute this file when loaded */
{
    tab_size.default = 5; /* set default value */
    load_commands("mycmds"); /* load my new cmds */
    do_save_state("epsilon"); /* save these changes */
}
```

For comparison, here are the lines you could add to your `einit.ecm` file instead, to make similar customizations without using EEL:

```
(set-variable tab-size 5)
(load-eel-from-path "mycmds.e" 2)
```

This technique, using an `einit.ecm` file as shown on page 164, is simpler than using a `changes.e` file, and doesn't require running the EEL compiler explicitly, so it's better for all but the most complex customizations.

Once you've learned a little EEL, you may want to modify some of Epsilon's built-in commands. We recommend that you keep your modifications to Epsilon in files other than the standard distributed source files. That way, when you get an update of Epsilon, you will find it easy to recompile your changes without accidentally loading in old versions of some of the standard functions.

You cannot redefine a function during that function's execution. Thus, changing the `load-bytes` command, for example, would seem to require writing a different command with the same functionality, and using each to load a new version of the other. You don't have to do this, however. Using the `-b` flag, you can load an entire system into Epsilon from bytecode files, not reading a state file at all. Epsilon does not execute any EEL functions while loading commands with the `-b` flag, so you can redefine any function using this technique.

To use this technique, first compile all the files that make up Epsilon. If you have a "make" utility program, you can use the makefile included with Epsilon to do this. Then start Epsilon with the `-b` flag. This loads the single bytecode file `epsilon.b`, which automatically loads all the others. The makefile then has Epsilon write a new state file using these definitions. If you have made extensive changes to Epsilon's commands, this method may be most convenient.

Chapter 8

Introduction to EEL



8.1 Epsilon Extension Language

The Epsilon Extension Language (EEL) allows you to write your own commands and greatly modify and customize the editor to suit your style. EEL provides a great deal of power. We used it to write all of Epsilon's commands. You can use it to write new commands, or to modify the ones that we provide.

We call EEL an *extension language* because you use it to extend the editor. Some people call such things *macro languages*. We use the term “macro” to refer to the keyboard macros you can create in Epsilon, or to EEL's C-like textual macros, but not to the commands or extensions you write in EEL.

EEL has these features:

- Block structure, with a syntax resembling the **C programming language**.
- The usual flow control: **if**, **while**, **for**, **do**, **switch** and **goto**. Additionally, EEL has a non-local goto facility provided by **setjmp** and **longjmp**.
- Complete set of data types, including **integers**, **arrays**, **structures**, and **pointers**. In addition, you may define new data types and allocate data objects dynamically.
- Subroutines with **parameter passing**. You may invoke subroutines **recursively**, and can designate any subroutine a command.
- Rich set of **arithmetic** and **logical** operators. EEL has all the operators of the C programming language.
- A powerful set of primitives. We wrote **all** of Epsilon's commands in EEL.
- Global variables accessible everywhere, and local variables accessible only in the current routine. EEL also has **buffer-specific variables** that change from buffer to buffer, and **window-specific variables** that have a different value in each window.

In addition, the runtime system provides a *simple debugger*, and an *execution profiler*.

Epsilon's source subdirectory contains the EEL source code to all Epsilon's commands. You may find it helpful to look at this source code when learning the extension language. Even after you've become a proficient EEL programmer, you probably will find yourself referring to the source code when writing your own extensions, to see how a particular command accomplishes some task.

8.2 EEL Tutorial

This section will take you step by step through the process of creating a new command using EEL. You will learn how to use the EEL compiler, a few control structures and data types, and a few primitive operations. Most important, this section will teach you the mechanics of writing extensions in EEL.

As our example, we will write a simplified version of the insert-file command called simple-insert-file. It will ask for the name of a file, and insert the contents of the file before point in the current buffer. We will write it a few lines at a time, each time having the command do more until the whole command works. When you write EEL routines, you may find this the way to go. This method allows you to debug small sections of code.

Start Epsilon in a directory where you want to create the files for this tutorial. Using the find-file command (Ctrl-x Ctrl-f), create a file with the name “learn.e”.

To write an extension, you: **write** the source code, **compile** the source code, **load** the compiled code, then **run** the command.

First, we write the source code. Type the following into the buffer and save it:

```

#include "eel.h"          /* standard definitions */

command simple_insert_file()
{
    char inserted_file[FNAMELEN];

    get_file(inserted_file, "Insert file", "");
    say("You typed file name %s", inserted_file);
}

```

Let's look at what the source code says. The first line includes the text of the file "eel.h" into this program, as though you had typed it yourself at that point.

Comments go between `/*` and `*/`.

The file "eel.h" defines some system-wide constants, and a few global variables. Always include it at the beginning of your extension files.

The line

```
command simple_insert_file()
```

says to define a command with the name `simple_insert_file`. The empty parentheses mean that this function takes no parameters. The left brace on the next line and the right brace at the end of the file delimit the text of the command.

Each command or subroutine begins with a sequence of local variable declarations. Our command has one, the line

```
char inserted_file[FNAMELEN];
```

which declares an array of characters called `inserted_file`. The array has a length of `FNAMELEN`. The constant `FNAMELEN` (defined in `eel.h`) may vary from one version to another. It specifies the maximum file name length, including the directory name. The semicolon at the end of the line terminates the declaration.

The next statement

```
get_file(inserted_file, "Insert file", "");
```

calls the built-in subroutine `get_file()`. This primitive takes three parameters: a character array to store the user's typed-in file name, a string with which to prompt the user, and a value to offer as a default. In this case, the Epsilon will prompt the user with the text between the double quotes (with a colon stuck on the end). We call a sequence of characters between double quotes a *string constant*.

When the user invokes this command, the prompt string appears in the echo area. Epsilon then waits for the user to enter a string, which it copies to the character array. While typing in the file name, the user may use Epsilon's file name completion and querying facility. This routine returns when the user hits the (Enter) key.

The next statement,

```
say("You typed file name %s", inserted_file);
```

prints in the echo area what file name the user typed in. The primitive `say()` takes one or more arguments. The first argument acts as a template, specifying what to print out. The `"%s"` in the above format string says to interpret the next argument as a character array (or a string), and to print that instead of the `"%s"`. In this case, for the second argument we provided `inserted_file`, which holds the name of the file obtained in the previous statement.

For example, say the user types the file name “foo.bar”, followed by `<Enter>`. The character array `inserted_file` would have the characters “foo.bar” in it when the `get_file()` primitive returns. Then the second statement would print out

```
You typed file name foo.bar
```

in the echo area.

One way to get this command into Epsilon is to run the EEL compiler to compile the source code into a form Epsilon can interpret, called a bytecode file. EEL source files end in “.e”, and the compiler generates a file of compiled binary object code that ends in “.b”. After you do that, you can load the .b file using the `load-bytes` command.

But an easier way that combines these steps is to use Epsilon’s `compile-buffer` command on Alt-F3. This command invokes the EEL compiler, as if you typed

```
eel filename
```

where *filename* is the name of the file you want to compile, and then (if there are no errors) loads the resulting bytecode file. You should get the message “learn.b compiled and loaded.” in the echo area.

Now that you’ve compiled and loaded `learn.b`, Epsilon knows about a command named `simple-insert-file`. Epsilon translates the underscores of command names to hyphens, so as to avoid conflicts with the arithmetic minus sign in the source text. So the name `simple_insert_file` in the eel source code defines `simple-insert-file` at command level.

Go ahead and invoke the command `simple-insert-file`. The prompt

```
Insert file:
```

appears in the echo area. Type in a file name now. You can use all Epsilon’s completion and querying facilities. If you press ‘?’, you will get a list of all the files. If you type “foo?”, you will get a list of all the files that start with “foo”. `<Esc>` and `<Space>` completion work. You can abort the command with Ctrl-g.

After you type a file name, this version of the command simply displays what you typed in the echo area.

Let’s continue with the `simple-insert-file` command. We will create an empty temporary buffer, read the file into that buffer, transfer the characters to our buffer, then delete the temporary buffer. Also, let’s get rid of the line that displays what you just typed. Make the file `learn.e` look like this:

```
#include "eel.h"          /* standard definitions */

command simple_insert_file()
{
    char inserted_file[FNAMELEN];
    char *original_buffer = bufname;

    get_file(inserted_file, "Insert file", "");
    zap("tempbuf");        /* make an empty buffer */
    bufname = "tempbuf"; /* use that buffer */
    if (file_read(inserted_file, 1) != 0)
        error("Read error: %s", inserted_file);
        /* copy the characters */
    xfer(original_buffer, 0, size());
        /* move back to buffer */
    bufname = original_buffer;
```

```

        delete_buffer("tempbuf");
    }

```

This version has one more declaration at the beginning of the command, namely

```
char *original_buffer = bufname;
```

This declares `original_buffer` to point to a character array, and initializes it to point to the array named `bufname`.

The value of the variable `bufname` changes each time the current buffer changes. For this reason, we refer to such variables as *buffer-specific variables*. At any given time, `bufname` contains the name of the current buffer. So this initialization in effect stores the name of the current buffer in the local variable `original_buffer`.

After the `get_file()` call, we create a new empty buffer named “tempbuf” with the statement “`zap("tempbuf");`”. We then make “tempbuf” the current buffer by setting the `bufname` variable with the following.

```
bufname = "tempbuf";
```

Now we can read the file in:

```

if (file_read(inserted_file, 1) > 0)
    error("Read error: %s", inserted_file);

```

This does several things. First, it calls the `file_read()` primitive, which reads a file into the current buffer. It returns 0 if everything goes ok. If the file doesn’t exist, or some other error occurs, it returns a nonzero error code. The actual return value in that case indicates the specific problem. This statement, then, executes the line

```
error("Read error: %s", inserted_file);
```

if an error occurred while reading the file. Otherwise, we move on to the next statement. The primitive `error()` takes the same arguments that `say()` takes. It prints out the message in the echo area, aborts the command, and drops any characters you may have typed ahead.

Now we have the text of the file we want to insert in a buffer named `tempbuf`. The next statement,

```
xfer(original_buffer, 0, size());
```

calls the primitive `xfer()`, which transfers characters from one buffer to another. The first argument specifies the name of the buffer to transfer characters to. The second and third arguments give the region of the current buffer to transfer. In this case, we want to transfer characters to `original_buffer`, which holds the name of the buffer from which we invoked this command. We want to transfer the whole thing, so we give it the parameters 0 and `size()`. The primitive `size()` returns the number of characters in the current buffer.

The last two statements return us to our original buffer and delete the temporary buffer.

The final version of this command adds several more details.

On the first line, we’ve added on `cx_tab['i']`. This tells Epsilon to bind the command to Ctrl-x i. We’ve added a new character pointer variable named `buf`, because we will use Epsilon’s `temp_buf()` subroutine for our temporary buffer rather than the wired-in name of “tempbuf”. This subroutine makes up an unused buffer name and creates it for us. It returns the name of the buffer.

```

#include "eel.h"          /* standard definitions */

char region_file[FNAMELEN];

command simple_insert_file() on cx_tab['i']
{
    char inserted_file[FNAMELEN], *buf;
    char *original_buffer = bufname;
    int err;

    iter = 0;
    get_file(inserted_file, "Insert file", region_file);
    mark = point;
    bufname = buf = temp_buf();
    err = file_read(inserted_file, 1);
    if (!err)
        xfer(original_buffer, 0, size());
    bufname = original_buffer;
    delete_buffer(buf);
    if (err)
        file_error(err, inserted_file, "read error");
    else
        strcpy(region_file, inserted_file);
}

```

Figure 8.1: The final version of simple-insert-file

The line

```
mark = point;
```

causes Epsilon to leave the region set around the inserted text. The `xfer()` will insert its text between mark and point. We've added the line `iter = 0;` to make the command ignore any numeric argument. Without this line, it would ask you for a file to insert over and over, if you accidentally gave it a numeric argument.

We now save the error code that `file_read()` returns so we can delete the temporary buffer in the event of an error. We also use the `file_error()` primitive rather than `error()` because the former will translate system error codes to text.

Finally, we added the line

```
char region_file[FNAMELEN];
```

to provide a default if you should execute the command more than once. Because this definition occurs outside of a function definition, the variable persists even after the command finishes. Variables defined within a function definition (local variables) go away when the function finishes. We copy the file name to `region_file` each time you use the command, and pass it to `get_file()` to provide a default value.

Chapter 9

Epsilon Extension Language



This chapter describes the syntax and semantics of EEL, the Epsilon Extension Language. Starting on page 439, we describe the built-in functions and variables (called *primitives*) of EEL. The tutorial that explains how to compile and load commands into Epsilon begins on page 391. You will find EEL very similar to the C programming language. A list of differences between EEL and C appears on page 429.

9.1 EEL Command Line Flags

To invoke the EEL compiler, type `eel filename`. If you omit the file name, the compiler will display a message showing its command line options.

Before the *filename*, you can optionally specify one or more command line switches. The EEL compiler looks for an environment variable named EEL before examining its command line, then “types in” the contents of that variable before the compiler’s real command line. Under Windows, the EEL compiler uses a registry entry named EEL (a “configuration variable”, as described on page 10), not an environment variable.

The EEL compiler has the following flags:

- b By default, in each file it compiles, the EEL compiler includes a variable definition that provides the full path to the original source file. It has a name such as `_loaded_eel_file_xyz` for a file `xyz.e`. This flag omits that variable; it’s used for compiling all EEL files that are part of Epsilon’s standard distribution. Epsilon’s `list-customizations` command uses the information provided by such variables.
- d*mac!def* This flag defines the textual macro *mac*, giving it the definition *def*, as if you had defined it using the `#define` command. The syntax `-dmac` defines the macro *mac*, giving it the definition (1). You can also use the syntax `-dmac=def`, but beware: if you run EEL via a .BAT or .CMD file, the system will replace any =’s with spaces, and EEL will not correctly interpret the flag.
- e This flag tells the compiler to exclude definitions from `#included` files when it writes the bytecode file. This results in smaller bytecode files. You can safely use this flag when compiling EEL files other than `epsilon.e` that only include the file `eel.h`, but it’s most useful with autoloading files. Epsilon will signal an error if you call a function using a variable whose definition has been omitted by `-e` in all loaded bytecode files.
- f This flag makes the compiler act as a filter, reading EEL code from `stdin` instead of a file, and writing its binary output to `stdout`. A file name on the command line is still required, but it is used only for error messages and debugging information.
- F This flag makes the compiler write its binary output to `stdout` instead of a bytecode file.
- i*directory* This flag sets the directories to search for files included with the preprocessor `#include` command. Precede each search directory with `-i`. If you use several `-i` flags on the command line, the compiler will search the directories in the order they appear. The compiler searches for included files using the following rules. First, if you use the syntax `#include "file.h"`, not `#include <file.h>`, EEL searches in the directory containing the current source file. The `-w1` flag makes it skip this step. Next, EEL searches in each directory specified by `-i`. If EEL still hasn’t found the include file, the `-w2` flag makes EEL give up at this point. Otherwise, EEL searches the `EPSPATH` configuration variable, looking for an `include` subdirectory of each directory. If there is no `EPSPATH` configuration variable, EEL searches a default `EPSPATH` (see page 12). When constructing this default `EPSPATH`, the `-w8` flag makes EEL omit any directory chosen based on the EEL compiler’s location.
- i- Makes the EEL compiler ignore all prior `-i` flags. This is useful if you use a configuration variable to always provide certain `-i` flags.

- n Makes the EEL compiler skip displaying its copyright message.
- o*file* Sets the output file. Normally EEL constructs the file name for the bytecode file based on the input file, with the .e extension replaced by “.b”, and puts the bytecode file in the current directory.
- p Makes the compiler display a preprocessed version of the file.
- q Suppress warning messages about unused local variables and function parameters.
- s Leave out debugging information from the bytecode file. Such a file takes up less space, and runs a bit faster. If you use this switch, though, you cannot use the debugger on this file, and the debug key Ctrl-(Scroll Lock) (except under Windows and Unix) will not work while such a function executes. We compiled the standard system with the -s flag. You may wish to recompile some files without this flag so you can trace through functions and see how they work.
- v Prints a hash mark each time the compiler encounters a function or global variable definition. Use it to follow the progress of the compiler.
- wNUM Bits in NUM control how EEL searches for include files. The 1 bit tells EEL to treat `#include "file.h"` the same as `#include <file.h>` and skip looking in the current source file's directory. The 2 bit tells EEL not to search for included files based on the EPSPATH. The 8 bit tells EEL that when constructing a default EPSPATH, it shouldn't consider the location of the EEL executable. See the description of the -i flag above. Values in the -w flag are cumulative, so -w1 -w2 is the same as -w3. Omit the number (use -w) to clear all bits.

An example using these switches is:

```
eel -s -p -v -dCODE=3 -oout -i/headers source >preproc
```

9.2 The EEL Preprocessor

EEL includes a preprocessor that can do macro substitution on the source text, among other things. You give preprocessor commands by including lines that start with “#” in your source text. A backslash character “\” at the end of a line makes the preprocessor command continue to the next line. This section lists the available preprocessor commands.

#define *identifier replacement-text*

This command defines a textual macro named *identifier*. When this identifier appears again in normal text (not in quotes), it is immediately replaced with the characters in the replacement text.

The rules for legal macro names are the same as the rules for identifiers in the rest of EEL: a letter or the underscore character “_”, followed by any number of letters, digits, or underscore characters. Identifiers which differ by case are different identifiers, so `mabel`, `maBel`, and `MABEL` could be three different macros. For clarity, it's best to use all upper case names for macros, and avoid such names otherwise.

When the EEL compiler starts, the macro `_EEL_` is predefined, with replacement text (1). The macros `UNICODE` and `BUNICODE` are also defined, with the same values. You can test for `UNICODE` to write EEL code that must also compile in older versions of Epsilon without Unicode support.

Note that these textual EEL macros are not related to keyboard macros. Only the EEL compiler knows about textual macros; Epsilon has no knowledge of them. You cannot bind a textual macro to a key, for example. Keyboard macros can be bound to a key, and the EEL compiler doesn't know anything about them, only the main Epsilon program. To further confuse matters, other editors refer to their extension languages

as macro languages, and call all editor extensions “macros”. In this manual, we never use the word “macro” to mean an editor extension written in EEL.

```
#define identifier(arg1,arg2,arg3,...) replacement-text
```

A macro with arguments is like a normal macro, but instances of the identifier in normal text must be followed by the same number of text sections (separated by commas) as there are arguments. Commas inside quotes or parentheses don’t separate text sections. Each of these text sections replace the corresponding identifier within the replacement text. For example, the preprocessor changes

```
#define COLOR(fg, bg) ((fg) + ((bg) << 4))
int modecol=COLOR(8, 3);
int mcol=COLOR(new_col(6,2),name_to_col("green"));
```

to

```
int modecol=((8) + ((3) << 4))
int mcol=((new_col(6,2))+((name_to_col("green"))<<4))
```

The command

```
#undef identifier
```

removes the effect of a prior #define for the rest of a compilation.

The command

```
#include <filename>
```

inserts the text in another file at this point in the source text. #include’s may be nested. In the above format, the EEL compiler searches for the file in each of the #include directories specified on the command line, or in a default location if none were specified. See page 397.

If you use quote marks (" ") instead of angle brackets (< >) around the file name of the #include command, the EEL compiler will first look in the directory of the original file for the included file, before searching the #include directories as above. With either delimiter, the compiler will ignore attempts to include a single file more than once in a compilation.

The command

```
#tryinclude <filename>
```

is the same as the #include command, except it’s not an error if EEL cannot locate the specified file. EEL just ignores the command in that case.

The EEL compiler keeps track of the current source file name and line number to provide error messages during compilation, and passes this information along in the bytecode file (unless you used the -s command line option to suppress this). Epsilon then uses this information for the EEL debugger and profiler, and displays it when certain errors occur. You can change the compiler’s notion of the current line and source file with the command

```
#line number "filename"
```

This makes the compiler believe the current file is *filename*, and the #line command appears on line *number* of it. If the file name is omitted, only the line number is changed.

```
#if constant-expression
    . . . text . . .
#endif
```

The `#if` command permits sections of the source text to be conditionally included. A constant expression (defined on page 423) follows the `#if`. If the value of the constant expression is nonzero, text from this point to a matching `#endif` command is included. Otherwise, that region is ignored. As part of the constant expression, the `defined()` keyword may be used; `defined(XYZ)` evaluates to 1 if a macro named XYZ has been defined, otherwise 0.

```
#if constant-expression
    . . . text . . .
#else
    . . . text . . .
#endif
```

If an `#else` command appears between the `#if` and the `#endif`, the text following the `#else` is ignored whenever the text preceding it is not. In other words, the text following the `#else` is ignored if the constant is nonzero.

```
#if constant-expression
    . . . text . . .
#elif constant-expression
    . . . text . . .
#else
    . . . text . . .
#endif
```

There may also be one or more `#elif` commands before an `#else` or `#endif` command. EEL evaluates each constant expression in turn, and includes the text following the first of these constant expressions that yields a nonzero value, skipping over all remaining commands in that block.

```
#ifdef identifier
    . . . text . . .
#endif

#ifndef identifier
    . . . text . . .
#endif
```

You can use the `#ifdef` command in place of the `#if` command. It ignores text between the command and a matching `#endif` if the identifier is not currently defined as a textual macro with the `#define` command. The text is included if the macro is defined. The `#ifndef` command is the same, but with the condition reversed. It includes the text only if the macro is undefined. Both commands may have `#else` or `#elif` sections, as with `#if`.

9.3 Lexical Rules

Comments in EEL begin with the characters `/*`, outside of any quotes. They end with the characters `*/`. The sequence `/*` has no effect while inside a comment, nor do preprocessor control lines.

You can also begin a comment with the characters `//`, outside of quotes. This kind of comment continues until the end of the line.

9.3.1 Identifiers

Identifiers in EEL consist of a letter or the underscore character “`_`”, followed by any number of letters, digits, or underscore characters. Upper case and lower case characters are distinct to the compiler, so `Ab` and `ab` are different identifiers. When you load an identifier into Epsilon, Epsilon converts underscores “`_`” to hyphens “`-`” and converts identifiers to lower case. For example, when invoking a command that has been defined in an EEL source file as `this_command()`, you type **this-command**. All characters are significant, and no identifier (or any token, for that matter) may be longer than 1999 characters.

The following identifiers are keywords, and you cannot use them for any other purpose:

<code>if</code>	<code>switch</code>	<code>struct</code>	<code>static</code>
<code>else</code>	<code>case</code>	<code>union</code>	<code>unsigned</code>
<code>for</code>	<code>default</code>	<code>keytable</code>	<code>enum</code>
<code>do</code>	<code>goto</code>	<code>typedef</code>	<code>color_class</code>
<code>while</code>	<code>sizeof</code>	<code>buffer</code>	<code>save_spot</code>
<code>return</code>	<code>char</code>	<code>window</code>	<code>save_var</code>
<code>break</code>	<code>short</code>	<code>command</code>	<code>spot</code>
<code>continue</code>	<code>int</code>	<code>on</code>	<code>on_exit</code>
<code>user</code>	<code>volatile</code>	<code>zeroed</code>	<code>color_scheme</code>
<code>byte</code>	<code>int32</code>		

The keywords `enum`, `unsigned`, and `static` have no function in the current version of EEL, but we reserve them for future use.

9.3.2 Numeric Constants

The term *numeric constant* collectively refers to decimal constants, octal constants, binary constants and hex constants.

A sequence of digits is a *decimal constant* (even if it begins with the digit 0, unlike the C language). The characters `0x` followed by a hexadecimal number are also recognized (the digits 0–9 and the letters `a–f` or the letters `A–F` form hexadecimal numbers). These are the *hex constants*. The characters `0b` followed by a binary number form a *binary constant*. A binary number contains only the digits 0 and 1. If it begins with a `0o`, it is an *octal constant* (base 8). Constants may contain `_` characters for readability, as in `1_000_000`; these are ignored.

All numeric constants in EEL are of type `int`.

9.3.3 Character Constants

Text enclosed in single quotes as in `'a'` is a *character constant*. The type of a character constant is `int`. Its value is the ASCII code for the character. Instead of a single character, an escape sequence can appear

between the quotes. Each escape sequence begins with a backslash, followed by a character from the following table. A backslash followed by any other character represents that character.

The special escape sequences are:

<code>\n</code>	newline character, <code>^J</code>
<code>\b</code>	backspace character, <code>^H</code>
<code>\t</code>	tab character, <code>^I</code>
<code>\r</code>	return character, <code>^M</code>
<code>\f</code>	form feed character, <code>^L</code>
<code>\yyy</code>	character with ASCII code <code>yyy</code> octal
<code>\xhh</code>	character with ASCII code <code>hh</code> hexadecimal
<code>\uhhhh</code>	character with code <code>hhhh</code> hexadecimal
<code>\u{h}</code>	character with code <code>h</code> hexadecimal
<code>\u[name]</code>	character with given Unicode <i>name</i>

For example, `'\''` represents the `'` character, `'\\'` represents the `\` character, `'\0'` represents the null character, and `'\n'`, `'\12'`, and `'\x0A'`, all represent the newline character (whose ASCII code is 12 in octal notation, base 8, and 0A in hexadecimal, base 16).

The `\u` sequence is followed by four hex digits, while the `\x` sequence is followed by only two, and so can only represent low-numbered characters. Enclose the hex digits in curly braces to directly mark the end of the hex number; one to four hex digits can appear within. The square bracket syntax recognizes any standard Unicode character name, such as `Greek Capital Letter Alpha`. These sequences are all equivalent: `\u[Greek Capital Letter Alpha]`, `\u0391`, `\u{391}`.

Anywhere a numeric constant is permitted, so is a character constant, and vice versa.

9.3.4 String Constants

Text enclosed in double quote characters (such as `"example"`) is a *string constant*. It produces a block of storage whose type is *array of char*, and whose value is the sequence of characters between the double quotes, with a null character (ASCII code 0) automatically added at the end. All the escape sequences for character constants work here too.

The compiler merges a series of adjacent string constants into a single string constant (before automatically adding a null character at the end). For example, `"sample" "text"` produces the same single block of storage as `"sampletext"`.

If an EEL file begins with a UTF-8 signature (a “byte order mark”), then the compiler interprets UTF-8 sequences in character and string constants. This method lets you include Unicode characters directly instead of using escape sequences. Select the “UTF-8” encoding via the `set-encoding` command to create such files. If a file does not begin with a UTF-8 signature, the compiler interprets bytes literally, for compatibility with older EEL code.

9.4 Scope of Variables

Variables may have two different kinds of “lifetimes”, or *scopes*. If you declare a variable outside of any function declaration, it is a *global variable*. If you declare it inside a function declaration, it is a *local variable*.

A local variable only exists while the function it is local to (the one you declared it in) is executing. It vanishes when the function returns, and reappears (with some different value) when the function executes

later. If you call the function recursively, each call of the function has its own value for the local variable. You may also declare a variable to be local to a block, in which case it exists only while code inside the block is executing. A local variable so declared only has meaning inside the function or block it is local to.

A global variable exists independently of any function. Any function may use it. If functions declared in different source files use the same global variable, the variable must be declared in both source files (or in files `#included` by both files) before its first use. If the two files have different initializations for the variable, only the first initialization has effect.

If a local variable has the same name as a global variable, the local masks the global variable. All references in the block to a variable of that name, from the local variable's definition until the end of the block it is defined in, are to the local variable. After the end of the block, the name again refers to the global variable.

You can declare any global variable to be *buffer-specific* using the `buffer` keyword. A buffer-specific variable has a value for each buffer and a default value. The default value is the value the variable has when you create a new buffer (and hence a new occurrence of the buffer-specific variable). When you refer to a buffer-specific variable, you normally refer to the part that changes from buffer to buffer. To refer to the default portion, append `".default"` to the variable name. For example, suppose the variable `foo` is buffer-specific. References to `foo` would then refer to the value associated with the current buffer. To refer to the default value, you would use the expression `foo.default`. (The syntax of appending `".default"` is available only when writing EEL programs, not when specifying a variable name to `set-variable`, for example.) When you save Epsilon's state using the `write-state` command, Epsilon saves only the default value of each buffer variable, not the value for the current buffer.

Global variables may also be declared *window-specific* using the `window` keyword. A window-specific variable has a separate value for each window and a default value. When Epsilon starts from a state file, it uses the default value saved in the state file to set up the first window. When you split a window, the new window's variables start off with the same values as the original window. Epsilon also uses the default value to initialize each new pop-up window. You can append `".default"` to refer to the default value of a window-specific variable.

9.5 Data Types

EEL supports a rich set of data types. First there are the *basic types*:

- `int` These are 64 bit signed quantities. These correspond to integers.
- `int32` These are 32 bit signed quantities. The value of an `int32` ranges from -2,147,483,648 to 2,147,483,647.
- `short` These are like ints, except they are only 16 bits. Thus the value ranges from -32768 to 32767.
- `char` These are 16 bit unsigned quantities. They correspond to characters. For example, the buffer primitive `curchar()` returns an object of type `char`. The values range from 0 to 65535.
- `byte` These are 8 bit unsigned quantities.
- `spot` These are references to buffer positions. A spot can remember a buffer position in such a way that after inserting or deleting characters in the buffer, the spot will still be between the same two characters. Like pointers, spots can also hold the special value zero. See page 442.

Besides basic types, there is an infinite set of types derived from these. They are defined recursively as follows:

pointer If t is some type, then *pointer to t* is also a type. Conceptually, this is the address of some object of type t . When you dereference an object of type *pointer to t* , the result is of type t .

array If t is some type, then *array of t* is also a type.

structure If t_1, \dots, t_n are types, then *structure of t_1, \dots, t_n* is also a type. Conceptually, a structure is a sequence of objects, where the j th object is of type t_j .

union If t_1, \dots, t_n are types, then *union of t_1, \dots, t_n* is also a type. Conceptually, a union is an object that can be of any of type t_1, \dots, t_n at different times.

function If t is a type, then *function returning t* is also a type.

Any function has a type, which is the type of the value it returns. If the function returns no value, it is of `int` type, but it is illegal to attempt to use the function's value.

Regardless of its type, you may declare any function to be a command (using the `command` keyword) if it takes no parameters. Commands like `named-command` on Alt-x will then complete on its name, but there is no other difference between commands and *subroutines* (user-defined functions which are not commands). Functions that the user is expected to invoke directly (by pressing a key, for example) are generally commands, while functions that act as helpers to commands are generally subroutines. Nothing prevents an EEL function from calling a command directly, though, and the user can invoke any subroutine directly as well (providing that it takes no arguments). Though a command may not have arguments, it may return a value (which is ignored when the user directly invokes it).

While it's possible for one command to call another command, commands are usually written to be called in the context of Epsilon's main loop (see page 578), not from other commands. New Epsilon commands should generally call EEL primitives and subroutines, not other commands, except in the case of a new command that simply adds some functionality to an existing single command.

9.5.1 Declarations

Declarations in EEL associate a type with an identifier. The structure of EEL declarations mimics the recursive nature of EEL types.

A *declaration* is of the form:

declaration:

type-specifier ;

type-specifier declarator-list ;

declarator-list:

declarator

declarator , *declarator-list*

A *type specifier* names one of the basic types, a structure or union (described on page 407), or a typedef, a type abbreviation (described on page 409).

type-specifier:

`char`

`short`

`int`

`struct struct-or-union-specifier`

`union struct-or-union-specifier`

`spot`

typedef-name
typedef-name:
identifier

A *declarator*, on the other hand, specifies the relationship of the identifier being declared to the type named by the type specifier. If this is a recursive type, the relationship of the identifier's type to the basic type of the type specifier is indicated by the form of the declarator.

Declarators are of the following form:

declarator:
identifier
 (*declarator*)
 * *declarator*
declarator [*constant-expression*]
declarator []
declarator ()

If *D* is a declarator, then (*D*) is identical to *D*. Use parentheses to alter the binding of composed declarators. We discuss this more on page 409.

9.5.2 Simple Declarators

In the simplest case, the identifier being declared is of one of the basic types. For that, the declarator is simply the identifier being declared. For example, the declarations

```
int length;
char this_character;
short small_value;
```

declare the type of the identifier `length` to be `int`, the type of `this_character` to be `char`, and the type of `small_value` to be `short`.

If the relationship between the identifier and the type specified in the type specifier is more complex, so is the declarator. Each type of declarator in the following sections contains exactly one identifier, and that is the identifier being declared.

9.5.3 Pointer Declarators

Pointer declarators are used in conjunction with type specifiers to declare variables of type *pointer to t*, where *t* is some type. The form of a pointer declarator is

* *declarator*

Suppose *T* is a type specifier and *D* is a declarator, and the declaration “*T D*;” declares the identifier embedded in *D* to be of type “... *T*”. Then the declaration *T *D*; declares the identifier in *D* to be of type “... *pointer to T*”. Several examples illustrate the concept.

```
int l;
int *lptr;
int **ldblptr;
```

Clearly, the first declaration declares `l` to be of type `int`. The type specifier is `int` and the declarator is `l`.

The second line is a little more complicated. The type specifier is still `int`, but the declarator is `*lptr`. Using the rule above, we see that `lptr` is a pointer to an `int`. This is immediately clear from the above if you substitute “`int`” for T , and “`lptr`” for D .

Similarly, the third line declares `ldb1ptr` to be a pointer to a pointer to an `int`.

9.5.4 Array Declarators

Array declarators are used in conjunction with type specifiers to declare objects of type *array of t* , where t is some type. The form of an array declarator is

declarator [*constant-expression*]

but you may omit the constant expression if

- An *initialized* global variable of type “*array of . . .*” is being defined. (See page 410.) In this case, the first constant-expression may be omitted, and the size of the array will be calculated from the initializer.
- A *function argument* (sometimes called a formal parameter) of type “*array of . . .*” is being declared. Since the type of the argument will be changed to “*pointer to . . .*” (as described on page 429) the first constant-expression may be omitted.

The rules for constant expressions appear on page 423.

Suppose T is a type specifier and D is a declarator, and the declaration “ $T D$,” declares the identifier embedded in D to be of type “. . . T ”. Then the declaration $T(D)[]$; declares the identifier to be of type “. . . array of T ”.

As an example, consider:

```
int (one_dim) [35];
int ((two_dim) [35]) [44];
```

The first line declares the identifier `one_dim` to be of type *array of int*.

The second line declares `two_dim` to be *array of array of int*. Clearly, we can have arbitrary multi-dimensional arrays by declaring the arrays in this manner.

As another example, consider the following:

```
char (*arg);
char (*argptr) [5];
char *(argary[5]);
```

From the preceding section, we know that the first line declares `arg` to be a pointer to a `char`. From this section, we see that the second line declares `argptr` to be of type *pointer to array of char*.

Compare this to the third line, which declares `argary` to be of type *array of pointer to char*.

When you have mixed declarators as you have in this example, you sometimes can elide parentheses according to the precedence rules of declarators. See section 9.5.7 for these precedences.

9.5.5 Function Declarators

Function declarators are used in conjunction with type specifiers to declare variables of type *function returning t*, where *t* is some type. The form of a function declarator is

declarator ()

or

declarator (*ansi-argument-list*)

Again, suppose *T* is a type specifier and *D* is a declarator, and the declaration “*T D*,” declares the identifier embedded in *D* to be of type “. . . *T*”. Then the declaration *T* (*D*) () ; declares the identifier to be of type “. . . *function returning T*”.

Consider:

```
char (c)();
char *(fpc());
char (*pfc)(int count, char *msg);
```

The first line declares *c* to be of type *function returning char*. The second line declares *fpc* to be a *function returning pointer to char*. The third line declares *pfc* to be of type *pointer to function returning char*. The third example also declares that *pfc* requires two parameters and gives their types; the first two examples provide no information about their functions’ parameters.

9.5.6 Structure and Union Declarations

This section describes how to define variables of type *structure of t₁, . . . , t_n*, where *t₁, . . . , t_n* are each types. First, we give an informal description, with examples, of how structures are often declared. A more formal description with BNF diagrams follows.

There is a special type-specifier, called a *structure-or-union specifier*, that defines structure and union types. This type-specifier has several forms.

The simplest form is seen in the following example:

```
struct {
    int field1;
    char name[30];
    char *data;
}
```

The field names of the structure are the identifiers being declared within the curly braces. These declarations look like variable declarations, but instead of declaring variables, they declare *field names*. The type of a particular field is the type the identifier would have if the declaration were a variable declaration.

The example above refers to a structure with fields named *field1*, *name*, and *data*, with types *int*, *array of char*, and *pointer to char*, respectively.

Use the structure-or-union specifier like the other type-specifiers (*int*, *short*, *char*, and *spot*) in declarations. For example:

```

struct {
    int field1;
    char name[30];
    char *data;
} rec, *recptr, recary[4];

```

declares `rec` to be a structure variable, `recptr` to be a pointer to a structure, and `recary` to be an array of (4) structures.

The structure-or-union-specifier may contain a *tag*, which gives a short name for the entire structure. For example, the type-specifier in the following example:

```

struct recstruct {
    int field1;
    char name[30];
    char *data;
};

```

creates a new type, `struct recstruct`, that refers to the structure being defined. Given this structure tag, we may define our structure variables in the following manner:

```

struct recstruct rec, *recptr, recary[4];

```

Structure (or union) tags also let you create self-referential types. Consider the following:

```

struct list {
    int data;
    struct list *next;
};

struct list list1, list2;

```

This creates a structure type `list`, which has a `data` field that's an `int`, and a `next` field that is a pointer to a `list` structure. A structure may not contain an instance of itself, but may contain (as in this example) a pointer to an object of its type.

More formally, a structure-or-union-specifier has the following form:

```

struct-or-union-specifier:
    struct-or-union-tag
    struct-or-union-tag { member-list }
    { member-list }
struct-or-union-tag:
    identifier
member-list:
    type-specifier declarator-list ;
    type-specifier declarator-list ; member-list

```

A description of how to use structures and unions in expressions appears on page 423.

9.5.7 Complex Declarators

As some of the examples thus far have shown, you can compose (combine) declarators to yield arbitrarily complicated types, like *function returning pointer to an array of 10 chars*:

```
char (*foo())[10];
```

When composing declarators, function and array declarators have the same precedence. They each take precedence over pointer declarators. So the example we used in section 9.5.5:

```
char *(fpc());
```

could have been written more simply as

```
char *fpc();
```

The rule that EEL follows for declarations is that the identifier involved is to be declared so that an expression with the form of the declarator has the type of the type specifier. This implies that the grouping of operators in a declarator follows the same rules as the operators do in an expression.

There are a few restrictions on the combinations of declarators when functions are involved (and so on the combinations of types). Functions may not return arrays, structures, unions, or functions, but they may return pointers to any of these. Similarly, functions may not be members of structures, unions, or arrays, but pointers to functions may be.

9.5.8 Typedefs

typedef-definition:

```
typedef type-specifier declarator-list ;
```

You can use typedefs to provide convenient names for complicated types. Once you define it, use a typedef as a type specifier (like `int`) in any declaration. A typedef definition looks just like a variable definition, except that the keyword `typedef` appears before the type specifier. The name of the typedef being defined appears instead of the variable name, and the typedef has the same type the variable would have had.

Typedefs only serve as abbreviations. They always create types that could be made in some other way. A variable declared using a typedef is just the same as a variable declared using the full specification. For example:

```
typedef short *NAME_LIST;
NAME_LIST nl, narray[20];
```

is equivalent to

```
short *nl, *narray[20];
```

9.5.9 Type Names

EEL's `sizeof` operator and its casting operator specify particular types using *type names*. A type name looks like a declaration of a single variable, except that the variable name is missing (as is the semicolon at the end). For example, `int *` is a type name referring to a pointer to an `int`.

type-name:

type-specifier abstract-declarator
abstract-declarator:
 empty
 (*abstract-declarator*)
 * *abstract-declarator*
 abstract-declarator [*constant-expression*]
 abstract-declarator []
 abstract-declarator ()
 abstract-declarator (*ansi-argument-list*)

Note that you could interpret a type name like `int *()` in two ways: either as a function returning a pointer to an int (like `int *foo()`;) or as a pointer to an int (like `int *(foo)`;) . EEL rules out the latter by requiring that a parenthesized *abstract-declarator* be nonempty. Given this, the system is not ambiguous, and an identifier can appear in only one place in each type name to make a legal declaration.

The same precedence rules apply to type names as to normal declarators (or to expressions). For example, the type name `char *[10]` refers to an array of 10 pointers to characters, but `char (*)[10]` refers to a pointer to an array of 10 characters.

9.6 Initialization

Declarations for the formal parameters of functions work just as described above, but you can additionally provide local and global variables with a specific initial value.

local-variable-definition:
 type-specifier local-declarator-list ;
local-declarator-list:
 local-declarator
 local-declarator , *local-declarator-list*
local-declarator:
 declarator
 declarator = *expression*

You can initialize a local variable with any expression so long as the corresponding assignment would be permitted. Since you cannot assign to variables with types such as “*array of . . .*” and “*structure of . . .*”, you cannot initialize such local variables at compile time. Local variables (those defined within a block) have undefined initial values if no explicit initialization is present.

global-variable-definition:
 type-specifier global-declarator-list ;
 global-modifier-list global-declarator-list ;
 global-modifier-list type-specifier global-declarator-list ;
global-modifier-list:
 global-modifier
 global-modifier global-modifier-list
global-modifier:
 buffer
 window
 zeroed

```

    user
    volatile
global-declarator-list:
    global-declarator
    global-declarator , global-declarator-list
global-declarator:
    declarator
    declarator = string-constant
    declarator = initializer
initializer:
    constant-expression
    { initializer-list }
    { initializer-list , }
initializer-list:
    initializer
    initializer , initializer-list

```

You may initialize a global variable of type “array of characters” with a string constant. If you omit the length of the array in a declaration with such an initialization, it’s set to just contain the initializing string (including its terminating null character).

If no explicit initialization is specified, variables defined globally are set to zero. If you provide a partial initialization (for example, if you specify the first 5 characters in a 10 character array), the remainder of the variable is set to zero. Initializers for global variables must involve only constant expressions known at compile time, whereas initializers for local variables may involve arbitrary expressions (including function calls, for example).

When Epsilon loads a file defining an initialized global variable and the variable was already defined to have the same type, the initialization has no effect: the variable’s value remains the same. If the new declaration specifies a different type for the variable, however, the variable’s value is indeed changed. (Actually, Epsilon only compares the sizes of the variables. If you redefine an integer as a four character array, Epsilon won’t apply the new initialization.) For example, suppose you declare `foo` to be an int and initialize it to 5. If you later load a file which redeclares `foo` to be an int and initializes it to 7, the value of `foo` would remain 5. If instead you redeclare `foo` to be a char and reinitialize it to ‘C’, then the value would change, since the size of a char is different from the size of an int.

To tell Epsilon that it must reinitialize the variable each time it reads a definition, use the `volatile` keyword. Every time you load a bytecode file containing such a variable definition, Epsilon will set the variable according to its initialization.

If you declare a global variable that is a number, spot, or pointer, the initializer must be a constant expression. In fact, if the variable is a spot or pointer, you can only initialize it with the constant zero. For example:

```

int i=3;
char *name="harold";

```

initializes the int variable `i` to be 3, and the character pointer `name` to point to the first character in the string “harold”. The variable `name` must be a local variable. If it were global, then you could initialize it only to zero, which is equivalent to not initializing it at all (see above).

If you declare a global array, you can initialize each element of the array. The initializer in this case would be a sequence of constant expressions, separated by commas, with the whole thing enclosed in braces `{}`. Consider the following examples:

```
int ary1[4] = { 10, 20, 30, 40 };
int ary2[ ] = { 10, 20, 30, 40 };
int ary3[4] = { 10, 20 };
```

Here we have `ary1` declared to be an array of 4 ints. We initialize the first element in the array to 10, the second to 20, and so on. The declaration of `ary2` does the same thing. Notice that the square brackets in the declarator are empty. The EEL compiler can tell from the initializer that the size must be 4. The declaration of `ary3` specifies the size of the array, but only initializes the first two elements. The compiler initializes the remaining two elements to zero.

The initializers for global structures are similar. The items between the curly braces are a sequence of expressions, with each expression's type matching the type of the corresponding field name. For example, the declaration:

```
struct {
    int f1;
    char f2;
    short f3;
} var = { 33, 't', 22 };
```

declares the variable `var` to be a structure with fields `f1`, `f2`, and `f3`, with types *int*, *char*, and *short* respectively. The declaration initializes the `f1` to 33, the character field `f2` to `'t'`, and the short field `f3` to 22.

You cannot initialize either unions or local structures. Global pointers may only be initialized to zero (which is equivalent to not initializing them at all).

If you initialize an array or structure which has subarrays or substructures, simply recursively apply the rules for initialization. For example, consider the following:

```
struct {
    char c;
    int ary1[3];
} var = { 't', { 3, 4, 5} };
```

This declares `var` to be a structure containing a character and an array of 3 ints. It initializes the character to `'t'`, and the array of ints so that the first element is 3, the second 4, and the third 5.

9.7 Statements

EEL has all of the statements of the C programming language. You can precede a statement by a *label*, an identifier followed by a colon, which you can use with the `goto` statement to explicitly alter the flow of control. Except where noted below, statements are executed in order.

9.7.1 Expression Statement

expression;

The expression is simply evaluated. This is the form of function calls and assignments, and is the most common type of statement in EEL.

9.7.2 If Statement

```
if ( expression )
    statement
```

If the value of *expression* is not zero, *statement* executes. Otherwise control passes to the statement after the if statement.

```
if ( expression )
    statement1
else
    statement2
```

If the value of *expression* is not zero, *statement1* executes. If the value of *expression* is zero, control passes to *statement2*.

9.7.3 While, Do While, and For Statements

```
while ( expression )
    statement
```

In a while loop, the *expression* is evaluated. If nonzero, the *statement* executes, and the expression is evaluated again. This happens over and over until the expression's value is zero. If the expression is zero the first time it is evaluated, *statement* is not executed at all.

```
do
    statement
while ( expression );
```

A do while loop is just like a plain while loop, except the statement executes *before* the expression is evaluated. Thus, the statement will always be evaluated at least once.

```
for ( expression1; expression2; expression3 )
    statement
```

In a for loop, first *expression1* is evaluated. Then *expression2* is evaluated, and if it is zero EEL leaves the loop and begins executing instructions after *statement*. Otherwise the statement is executed, *expression3* is evaluated, and *expression2* is evaluated again, continuing until *expression2* is zero.

You can omit any of the expressions. If you omit *expression2*, it is like *expression2* is nonzero. while (*expression*) is the same as for (; *expression*;). The syntax for (;;) creates an endless loop that must be exited using the break statement (or one of the other statements described below).

9.7.4 Switch, Case, and Default Statements

```
switch ( expression )
    statement

case constant-expression: statement

default: statement
```

Statements within the *statement* following the switch (which is usually a block, as described below) are labeled with constant expressions using case. The *expression* is evaluated (it must yield an int), and

Epsilon branches to the case statement with the matching constant. If there is no match, Epsilon branches to the default statement if there is one, and skips over the switch statement if not.

A case or default statement associates with the smallest surrounding switch statement. Each switch statement must have at most one case statement with any given value, and at most one default statement.

9.7.5 Break and Continue Statements

```
break;
```

This statement exits from the smallest containing for, while, do while or switch statement. The break statement must be the last statement in each case if you don't want execution to "fall through" and execute the statements for the following cases too.

```
continue;
```

The continue statement immediately performs the test for the smallest enclosing for, while, or do while statement. It is the same as jumping to the end of the *statement* in each of their definitions. In the case of for, *expression3* will be evaluated first.

9.7.6 Return Statement

```
return;
```

```
return expression;
```

The return statement exits from the function it appears in. The first form returns no value, and produces an error message if you called the function in a way that requires a value. The second form returns *expression* as the value of the function. It must have the same type as you declared the function to be. It is not an error for the value to be unused by the caller.

If execution reaches the end of a function definition, it is the same as if return; were there.

9.7.7 Save_var and Save_spot Statements

```
statement:
```

```
    save_var save-list;
```

```
    save_spot save-list;
```

```
save-list:
```

```
    save-item
```

```
    save-item , save-list
```

```
save-item:
```

```
    identifier
```

```
    identifier = expression
```

```
    identifier modify-operator expression
```

```
    identifier ++
```

```
    identifier --
```

The save_var statement tells Epsilon to remember the current value of a variable, and set it back to its current value when the function that did the save_var exits. Epsilon will restore the value no matter how

the function exits, even if it calls another function which signals an error, and this aborts out of the calling function.

You can provide a new value for the variable at the same time you save the old one. Epsilon first saves the old value, then assigns the new one. You can use any of the assignment operators listed on page 421, as well as the ++ and -- operators.

For example, this command plays a note at 440 Hz for one second, without permanently changing the user's variable settings for the bell (in versions of Epsilon that support changing the bell's frequency and duration).

```
command play_note()
{
    save_var beep_frequency = 440;
    save_var beep_duration = 100;
    ding();          /* uses beep_ variables */
}
```

The `save_spot` statement functions like `save_var`, but it creates a *spot* (see page 442) in the current buffer to hold the old value. The spot will automatically go away when the function exits. Use `save_spot` instead of `save_var` when you wish to save a buffer position, and you want it to stay in the right place even if the buffer contents change.

The `save_var` and `save_spot` statements can apply to global variables with “simple” types: those that you can directly assign to with the = operator. They don't work on structures, for example, or on local variables.

Although the `save_var` and `save_spot` statements resemble variable declarations, they are true statements. You can use the `if` statement (above), for example, to only save a variable in certain cases. These statements operate with a “stack” of saved values, so that if you save the same variable twice in a function, only the first setting will have an effect on the final value of the variable. (Repeated `save` statements take up space on the saved value stack, however, so they should be avoided.) When you save a buffer-specific or window-specific variable, Epsilon remembers which buffer's or window's value was saved, and restores only that one.

The `restore_vars()` primitive restores all variables saved in the current function. After a `restore_vars()`, future modifications to any saved variables won't be undone.

9.7.8 On_exit Statement

statement:

`on_exit statement`

An `on_exit` statement tells Epsilon to run some code later, when the current function exits. It can be used to clean up temporary data. As with the `save_var` statement in the previous section, Epsilon will run the specified code no matter how the function exits, even if it calls another function which signals an error, and this aborts out of the calling function.

The statement to be executed later can use local or global variables, call other functions, and so forth. It can be a block or other complex type of statement, but it cannot use certain control flow statements: `switch`, `case`, `default`, `break`, `continue`, `return`, `goto`, or labels. And it cannot use `on_exit`, `save_var`, `save_spot`, or `restore_vars()` itself. (If the `on_exit` statement's code calls a function, that function can use any of these.)

The `restore_vars()` primitive also causes all pending `on_exit` statements to be executed at once, just as if the current function were about to exit.

All `on_exit`, `save_var`, and `save_spot` statements push their pending operations onto a stack, and they are executed in order when the function returns or when `restore_vars()` is used, newest to oldest.

9.7.9 Goto and Empty Statements

```
goto label;

label: statement
```

The next statement executed after the `goto` will be the one following the *label*. It must appear in the same function as the `goto`, but may be before or after.

```
;
```

This null statement is occasionally used in looping statements, where all the “work” of the loop is done by the expressions. For example, a loop that calls a function `foo()` repeatedly until it returns zero can be written as

```
while (foo()) ;.
```

9.7.10 Block

```
{
  declarations
  statements
}
```

Anywhere you can have a statement, you can have a *block*. A block contains any number of local variable *declarations* or *statements* (including zero). The variables declared in the block are local to the block, and you may only use them in the following *statements* (or in statements contained in those statements). A block’s declarations can be mixed in freely among its statements. The body of a function definition is itself a block.

9.8 Conversions

When a value of a certain type is changed to another type, a *conversion* occurs.

When a number of some type is converted to another type of number, if the number can be represented in the latter type its value will be unchanged. All possible characters can be represented as ints or short ints, and all short ints can be represented as ints, so these conversions yield unchanged values.

Technically, Epsilon will sign-extend a short int to convert it to an int, but will pad a character with zero bits on the left to convert it to an int or short int. Converting a number of some type to a number of a shorter type is always done by dropping bits.

A pointer may not be converted to an int, or vice versa, except for function pointers. The latter may be converted to a short int, or to any type that a short int may be converted to. A pointer to one type may be converted to a pointer to another type, as long as neither of them is a function pointer.

All operators that take numbers as operands will take any size numbers (characters, short ints, or ints). The operands will be converted to int if they aren’t already ints. Operators that yield numbers always produce ints.

Highest Precedence	
l-to-r	() [] -> .
r-to-l	All unary operators (see below)
l-to-r	* / %
l-to-r	+ -
l-to-r	<< >>
l-to-r	> < >= <=
l-to-r	== !=
l-to-r	&
l-to-r	^
l-to-r	
l-to-r	&&
l-to-r	
l-to-r	? :
r-to-l	All assignment operators (see below)
l-to-r	,
Lowest Precedence	
Assignment operators are: = *= /= %= += -=	
<<= >>= &= ^= =	
Unary operators are: * & - ! ~	
++ -- sizeof (type-name)	

Figure 9.1: Operator Precedence

9.9 Operator Grouping

In an expression like

10 *op1* 20 *op2* 30

the compiler determines the rules for grouping by the *precedence* and *associativity* of the operators *op1* and *op2*. Each operator in EEL has a certain precedence, with some precedences higher than others. If *op1* and *op2* have different precedences, the one with the higher precedence groups tighter. In table 9.1, operators with higher precedences appear on a line above operators with lower precedences. Operators with the same precedence appear on the same line.

For example, say *op1* is + and *op2* is *. Since *’s line appears above +’s, * has a higher precedence than + and the expression 10 + 20 * 30 is the same as 10 +(20 * 30).

If two operators have the same precedence, the compiler determines the grouping by their associativity, which is either left-to-right or right-to-left. All operators of the same precedence have the same associativity. For example, suppose *op1* is - and *op2* is +. These operators have the same precedence, and associate left-to-right. Thus 10 - 20 + 30 is the same (10 - 20) + 30. All operators on the same line in the table have the same precedence, and their associativity is given with either “l-to-r” or “r-to-l.”

Enclosing an expression in parentheses alters the grouping of operators. It does not change the value or type of an expression itself.

9.10 Order of Evaluation

Most operators do not guarantee a particular order of evaluation for their operands. If an operator does, we mention that fact in its description below. In the absence of such a guarantee, the compiler may rearrange calculations within a single expression as it wishes, if the result would be unchanged ignoring any possible side effects.

For example, if an expression assigns a value to a variable and uses the variable in the same expression, the result is undefined unless an operator that guarantees order of evaluation occurs at an appropriate point.

Note that parentheses do not alter the order of evaluation, but only serve to change the grouping of operators. Thus in the statement

```
i = foo() + (bar() + baz());
```

the three functions may be called in any order.

9.11 Expressions

9.11.1 Constants and Identifiers

expression:

numeric-constant

string-constant

identifier

color_class identifier

The most basic kinds of expressions are numeric and string constants. Numeric constants are of type “int”, and string constants are of type “array of character”. However, EEL changes any expression of type “array of . . .” into a pointer to the beginning of the array (of type “pointer to . . .”). Thus a string constant results in a pointer to its first character.

An identifier is a valid expression only if it has been previously declared as a variable or function. A variable of type “array of . . .” is changed to a pointer to the beginning of the array, as described above.

Some expressions are called *lvalue expressions*. Roughly, lvalue expressions are expressions that refer to a changeable location in memory. For example, if `foo` is an integer variable and `func()` is a function returning an integer, then `foo` is an lvalue, but `func()` is not. The `&` and `.` operators, the `++` and `--` operators, and all assignment operators require their operands to be lvalues. Only the `*`, `[]`, `->`, and `.` operands return lvalues.

An identifier which refers to a variable is an lvalue if its type is an integer, a spot, a pointer, a structure, or a union, but not if its type is an array or function.

If an identifier has not been previously declared, and appears in a function call as the name of the function, it is implicitly declared to be a function returning an int.

If the name of a previously declared function appears in an expression in any context other than as the function of a function call, its value is a function pointer to the named function. Function pointers may not point to primitive functions.

For example, if `foo` is previously undeclared, the statement `foo(1, 2);` declares it as a function returning an int. If the next statement is `return foo;`, a pointer to the function `foo()` will be returned.

Once a color class `newclass` has been declared, you can refer to it by using the special syntax `color_class newclass`. This provides a numeric code that refers to the particular color class. It's used in conjunction with the primitives `alter_color()`, `add_region()`, `set_character_color()`, and others. See page 113 for basic information on color classes, and page 425 for information on declaring color classes in EEL.

9.11.2 Unary Operators

expression:

```
! expression
* expression
& expression
- expression
~ expression
sizeof expression
sizeof( type-name )
( type-name ) expression
++ expression
-- expression
expression ++
expression --
```

The `!` operator yields one if its operand is zero, and zero otherwise. It can be applied to pointers, spots, or numbers, but its result is always an int.

The unary `*` operator takes a pointer and yields the object it points to. If its operand has type “pointer to . . .”, the result has type “. . .”, and is an lvalue. You can also apply `*` to an operand of type “spot”, and the result is a number (a buffer position).

The unary `&` operator takes an lvalue and returns a pointer to it. It is the inverse of the `*` operator, and its result has type “pointer to . . .” if its operand has type “. . .”. (You cannot construct a spot by applying the `&` operator to a position. Use the `alloc_spot()` primitive described on page 442.)

The unary `-` and `~` operators work only on numbers. The first negates the given number, and the second flips all its bits, changing ones to zeros and zeros to ones.

The `sizeof` operator yields the size in bytes of an object. You can specify the object as an expression or with a type name (described on page 409). In the latter case, `sizeof` returns the size in bytes of an object of that type. Characters and shorts require two bytes, and ints four bytes. An array of 10 ints requires 40 bytes, and this is the number `sizeof(int [10])` will give, not 10.

An expression with a parenthesized type name before it is a *cast*. The cast converts the expression to the named type using the rules beginning on page 416, and the result is of that type. Specify the type using a type name, described on page 409.

The `++` and `--` operators increment and decrement their lvalue operands. If the operator appears before its operand, the value of the expression is the new value of the operand. The expression `(++var)` is the same as `(var += 1)`, and `(--var)` is the same as `(var -= 1)`. You can apply these operators to pointers, in which case they work as described under pointer addition below.

If the `++` or `--` operators appear after their operand, the operand is changed in the same way, but the value of the expression is the value of the operand *before* the change. Thus the expression `var++` has the same value as `var`, but `var` has a different value when you reference it the next time.

9.11.3 Simple Binary Operators

expression:

```

expression + expression
expression - expression
expression * expression
expression / expression
expression % expression
expression == expression
expression != expression
expression < expression
expression > expression
expression <= expression
expression >= expression
expression && expression
expression || expression
expression & expression
expression | expression
expression ^ expression
expression << expression
expression >> expression

```

The binary `+` operator, when applied to numbers, yields the sum of the numbers. One of its operands may also be a pointer to an object in an array. In this case, the result is a pointer to the same array, offset by the number to another object in the array. For example, if `p` points to an object in an array, `p + 1` points to the next object in the array and `p - 1` points to the previous object, regardless of the object's type.

The binary `-` operator, when applied to numbers, yields the difference of the numbers. If the first operand is a pointer and the second is a number, the rules for addition of pointers and numbers apply. For example, if `p` is a pointer, `p - 3` is the same as `p + -3`.

Both operands may also be pointers to objects in the same array. In this case the result is the difference between them, measured in objects. For example, if `arr` is an array of ten ints, `p1` points to the third int, and `p2` points to the eighth, then `p1 - p2` yields the int `-5`. The result is undefined if the operands are pointers to different arrays.

The binary `*` operator is for multiplication, and the `/` operator is for division. The latter truncates toward 0 if its operands are positive, but the direction of truncation is undefined if either operand is negative. The `%` operator provides the remainder of the division of its operands, and `x % y` is always equal to `x - (x / y) * y`. All three operators take only numbers and yield ints.

The `==` operator yields one if its arguments are equal and zero otherwise. The arguments must either both be numbers, both spots, or both pointers to objects of the same type. However, if one argument is the constant zero, the other may be a spot or any type of pointer, and the expression yields one if the pointer is null, and zero otherwise. The `!=` operator is just like the `==` operator, but returns one where `==` would return zero, and zero where `==` would return one. The result of either operator is always an int.

The `<`, `>`, `<=`, and `>=` operators have a value of one when the first operand is less than, greater than, less than or equal to, or greater than or equal to (respectively) the second operand. The operands may both be numbers, they may be pointers to the same array, or one may be a pointer or spot and the other zero. In the last case, Epsilon returns values based on the convention that a null pointer or spot is equal to zero and a non-null one is greater than zero. The result is undefined if the operands are pointers to different arrays of

the same type, and it is an error if they are pointers to different types of objects, or if one is a spot and the other is neither a spot nor zero.

The `&&` operator yields one if both operands are nonzero, and zero otherwise. Each operand may be a pointer, spot, or number. Moreover, the first operand is evaluated first, and if it is zero, the second operand will not be evaluated. The result is an int.

The `||` operator yields one if either of its operands are nonzero, and zero if both are zero. Each operand may be a pointer, spot, or number. The first operand is evaluated first, and if it is nonzero, the second operand will not be evaluated. The result is an int.

The `&` operator yields the bitwise AND of its numeric operands. The `|` and `^` operators yields the bitwise OR and XOR (respectively) of their numeric operands. The result for all three is an int. A bit in the result of an AND is on if both corresponding bits in its operands are on. A bit in the result of an OR is on if either of the corresponding bits in its operands are on. A bit in the result of an XOR is on if one of the corresponding bits in its operands is on and the other is off.

The `<<` operator yields the first operand with its bits shifted to the left the number of times given by the right operand. The `>>` operator works similarly, but shifts to the right. The former fills with zero bits, and the latter fills with one bits if the first operand was negative, and zero bits otherwise. If the second operand is negative or greater than 31, the result is undefined. Both operands must be numbers, and the result is an int.

9.11.4 Assignment Operators

expression:

expression = *expression*

expression *modify-operator* *expression*

modify-operator:

`+=`

`-=`

`*=`

`/=`

`%=`

`&=`

`|=`

`^=`

`<<=`

`>>=`

The plain assignment operator `=` takes an lvalue (see page 418) as its first operand. The object referred to by the lvalue is given the value of the second operand. The types of the operands may both be numbers, spots, pointers to the same type of object, or compatible structures. If the first operand is a pointer or spot and the second is the constant zero, the pointer or spot is made null. The value of the expression is the new value of the first operand, and it has the same type.

The other kinds of assignment operators are often used simply as abbreviations. For example, if `a` is a variable, `a += (b)` is the same as `a = a + (b)`. However, the first operand of an assignment is only evaluated once, so if it has side effects, they will only occur once.

For example, suppose `a` is an array of integers with values 10, 20, 30, and so forth. Suppose `p()` is a function that will return a pointer to the first element of `a` the first time it's called, then a pointer to the second element, and so forth. After the statement `*p() += 3;`, `a` will contain 13, 20, 30. After `*p() = *p() + 3;`, however, `a` is certain not to contain 13, 20, 30, since `p()` will never return a pointer to the same

element of a twice. Because the order of evaluation is unspecified with these operators, the exact result of the latter statement is undefined (either 10, 13, 30 or 23, 20, 30).

The result of all these assignment statements is the new value of the first operand, and will have the same type. The special rules for mixing pointers and ints with the + and - operators also apply here.

9.11.5 Function Calls

expression:

expression ()

expression (*expression-list*)

expression-list:

expression

expression , *expression-list*

An expression followed by a parenthesized list of expressions (arguments) is a function call. Usually the first expression is the name of a function, but it can also be an expression yielding a function. (The only operator that yields a function is the unary * when applied to a function pointer.) The type of the result is the type of the returned value. If the function returns no value, the expression must appear in a place where its value is not used. You may call any function recursively.

If an identifier that has not been previously declared appears as the name of the function, it is implicitly declared to be a function returning an int.

Each argument is evaluated and a copy of its value is passed to the function. Character and short arguments are converted to ints in the process. Aside from this, the number and type of arguments must match the definition of the function. The order of evaluation of the arguments to a function is undefined.

Since only a copy of each parameter is passed to the function, a simple variable cannot be altered if its name only appears as the argument to a function. To alter a variable, pass a pointer to it, and have the function modify the object pointed to. Since an array is converted to a pointer whenever its name occurs, an array that is passed to a function can indeed be altered by the function. Numbers, spots, and pointers may be parameters, but structures, unions, or functions cannot be. Pointers to such things are allowed, of course.

An EEL function can call not just other EEL functions, but also any of Epsilon's built-in functions, known as primitives. These are listed in the next chapter. An EEL function can also call a keyboard macro as a function. The word "function" refers to any of the various types of routines that a command written in EEL can call. These include other commands or subroutines (themselves written in EEL), primitives that are built into Epsilon and cannot be changed, and keyboard macros (see page 155). Textual macros that are defined with the #define preprocessor statement are *not* functions.

Each function may require a certain number of arguments and may return a value of a particular type. Keyboard macros, however, never take arguments or return a value.

9.11.6 Miscellaneous Operators

expression:

expression ? *expression* : *expression*

expression , *expression*

expression [*expression*]

expression -> *identifier*

expression . *identifier*

The conditional operator `? :` has three operands. The first operand is always evaluated first. If nonzero, the second operand is evaluated, and that is the value of the result. Otherwise, the third operand is evaluated, and that is the value of the result. Exactly one of the second and third operands is evaluated. The first operand may be a number, spot, or pointer. The second and third operands may either both be numbers, both spots, both pointers to the same type of object, or one may be a pointer or spot and the other the constant zero. In the first case the result is an int, and in the last two cases the result is a spot or a pointer of the same type.

The `,` operator first evaluates its first argument and throws away the result. It then evaluates its second argument, and the result has that value and type. In any context where a comma has a special meaning (such as in a list of arguments), EEL assumes that any commas it finds are used for that special meaning.

The `[]` operator is EEL's subscripting operator. Because of the special way that addition of a pointer and a number works, we can define the subscripting operator in terms of other operators. The expression `e1[e2]` is the same as `*((e1)+(e2))`, and since addition is commutative, also the same as `e2[e1]`. In practice, subscripting works in the expected way. Note that the first object in an array has subscript 0, however. One of the operands must be a pointer and the other a number. The type of the result is that of the pointed-to object.

The `.` operator disassembles structures or unions. Its operand is an lvalue which is a structure or union. After the `.` an identifier naming one of the operand's members must appear. The result is an lvalue referring to that member.

The `->` operator is an abbreviation for a dereference (unary `*`) followed by a member selection as above. Its operand is a pointer to a structure or union, and it is followed by the name of one of the structure's or union's members. The result is an lvalue referring to that member. The expression `strptr->membername` is the same as the expression `(*strptr).membername`.

9.12 Constant Expressions

A constant expression is an expression which does not contain certain things. It may not have references to variables, string constants, or function calls. No subexpressions may have a type of spot, structure, union, array, or pointer. It may have numeric constants, character constants, and any operators that act on them, and the `sizeof` operator may appear with any operand.

Additionally, for constant expressions in preprocessor lines, you can test if a macro `m` has been defined by writing `defined(m)`. This expression evaluates to 1 if a macro by that name has been defined, 0 if not.

The term "the constant zero" means a constant expression whose value is zero, not necessarily a numeric constant.

9.13 Global Definitions

program:

global-definition

global-definition program

global-definition:

function-definition

global-variable-definition

keytable-definition

typedef-definition

color-class-definition

Each file of EEL code consists of a series of definitions for global variables and functions. Global variable definitions have the same format as local variable definitions. The first definition of a global variable Epsilon receives determines the initial value of the variable, and later initializations have no effect, unless you use the `volatile` keyword when defining the variable (see page 411). If the first definition provides no explicit initialization, the variable is filled with zeros or null pointers as appropriate, depending on its type.

You can declare any global variable (except a key table or color class) to be buffer-specific by placing the keyword `buffer` before the type specifier. When the definition is first read in, its initializer determines the value of the variable for each buffer that then exists, and also the default value of the variable. Whenever you create a new buffer (and hence a new copy of the buffer-specific variable), the variable's value in that buffer is set to the default value.

Similarly, you can declare any global variable except a key table or color class to be window-specific by placing the keyword `window` before the type specifier. When the definition is first read in, its initializer determines the value of the variable for each window that then exists, and also the default value of the variable. Whenever you split a window in two, the new window inherits its initial value for the window-specific variable from the original window. Epsilon uses the default value of a window-specific variable when it creates the first tiled window while starting up, and when it creates pop-up windows.

Epsilon's `write-state` command writes a new state file containing all variables, EEL functions, macros, colors, and so forth that Epsilon knows about. The file includes the current values of all numeric variables, all global character array variables, and any structures or unions containing just these types. But Epsilon doesn't save the values of variables containing pointers or spots, and sets these to zero as it writes a state file. You can put the `zeroed` keyword before the definition of a variable of any type to tell Epsilon to zero that variable when it writes a state file.

In commands like `set-variable`, Epsilon distinguishes between user variables and system variables, and only shows the former in its list of variables you can set. By default, each global variable you define is a system variable that users will not see. Put the `user` keyword before a variable's definition to make the variable a user variable.

9.13.1 Key Tables

keytable-definition:

`keytable keytable-list ;`

keytable-list:

identifier

identifier , keytable-list

A key table is a set of bindings, one for each key on the keyboard, with keys modified by control, alt, and shift counted as separate keys. Various mouse actions and system events are also represented by special key codes. Each entry in the key table contains a short integer, which is an index into the name table. In other words, each entry corresponds to a named Epsilon object, either a command, subroutine, keyboard macro, or another key table.

You can declare a key table by using the `keytable` keyword in place of the type specifier in a global variable definition. A key table definition can contain no initialization, just `keytable` followed by a list of comma-separated key table names and a semicolon. A key table acts like an array of short ints, but you can also use it in the on part of a function definition (as described below).

Key codes, the values that index a key table, can be very large numbers. Looping through all possible key codes with a simple `for (i = 0; i < MAXKEYS; i++)` statement is far too slow; see page 579 for the right way to write such loops.

9.13.2 Color Classes

```

color-class-definition:
    color_class color-class-list ;
    color_scheme color-scheme-list ;
color-class-list:
    color-class-item
    color-class-item , color-class-list
color-class-item:
    identifier
    identifier color_scheme string-constant = color-pair
    identifier { color-scheme-spec-list }
    identifier = color-pair
color-scheme-spec-list:
    color-scheme-spec
    color-scheme-spec color-scheme-spec-list
color-scheme-spec:
    color_scheme string-constant = color-pair ;
color-scheme-list:
    color-scheme-item
    color-scheme-item , color-scheme-list
color-scheme-item:
    string-constant
    string-constant color_class identifier = color-pair
    string-constant { color-class-spec-list }
color-class-spec-list:
    color-class-spec
    color-class-spec color-class-spec-list
color-class-spec:
    color_class identifier = color-pair ;
color-pair:
    color_class identifier
    constant-expression
    constant-expression on constant-expression

```

A color class specifies a particular pair of foreground and background colors Epsilon should use on a certain part of the screen, or when displaying a certain type of text. For example, Epsilon uses the color class `c_keyword` to display keywords in C-like languages. More precisely, the color class specifies which foreground/background pair of colors to display under each defined color scheme. If the user selects a different color scheme, Epsilon will immediately begin displaying C keywords using the `c_keyword` color pair defined for the new scheme.

Before you use a color class in an expression like `set_character_color(pos1, pos2, color_class c_keyword);`, you must declare the color class (outside of any function definition) using the `color_class` keyword:

```
color_class c_keyword;
```

When you declare a new color class, you may wish to specify the colors to use for a particular color scheme using the `color_scheme` keyword:

```
color_class c_keyword
    color_scheme "standard-gui" = black on white;
color_class c_keyword
    color_scheme "standard-color" = green on black;
```

If you have many color definitions all for the same color class, you can use this syntax:

```
color_class c_keyword {
    color_scheme "standard-gui" = black on white;
    color_scheme "standard-color" = green on black;
};
```

Similarly, if you have many color definitions for the same color scheme, you can avoid repeating it by writing:

```
color_scheme "standard-gui" {
    color_class c_keyword = black on white;
    color_class c_function = blue on white;
    color_class c_identifier = black on white;
};
```

To specify the particular foreground and background colors for a color class (using the syntax *foreground on background*), you can use these macros defined in `eel.h`:

```
#define black          MAKE_RGB(0, 0, 0)
#define dark_red       MAKE_RGB(128, 0, 0)
#define dark_green     MAKE_RGB(0, 128, 0)
#define brown          MAKE_RGB(128, 128, 0)
// etc.
```

See that file for the current list of named colors. These functions use the `MAKE_RGB()` macro, providing particular values for red, green, and blue. You can use this macro yourself, in a color class definition, to specify precise colors:

```
color_scheme "my-color-scheme" {
    color_class c_keyword = MAKE_RGB(223, 47, 192) on yellow;
};
```

There are several other macros useful with color expressions:

```
#define MAKE_RGB(rd,grn,bl) (((rd) << 16) + ((grn) << 8) + (bl))
#define GETRED(rgb)         (((rgb) >> 16) & 0xff)
#define GETGREEN(rgb)       (((rgb) >> 8) & 0xff)
#define GETBLUE(rgb)        ((rgb) & 0xff)
```

The `GETRED()`, `GETGREEN()`, and `GETBLUE()` macros take a color expression created with `MAKE_RGB()` and extract one of its three components, which are always numbers from 0 to 255.

The foreground color for a color class may also include font style bits, by or'ing any of the macros `EFONT_BOLD`, `EFONT_UNDERLINED`, and `EFONT_ITALIC` into the color code.

The `ETRANSSPARENT` macro is a special code that may be used in place of a background color. It tells Epsilon to substitute the background color of the "text" color class in the current color scheme. The following three examples are all equivalent:

```
color_class text color_scheme "standard-gui" = yellow on red;
color_class c_keyword color_scheme "standard-gui" = blue on red;

color_class text color_scheme "standard-gui" = yellow on red;
color_class c_keyword color_scheme "standard-gui" = blue
    on ETRANSPARENT;

color_class text color_scheme "standard-gui" = yellow on red;
color_class c_keyword color_scheme "standard-gui" = blue;
```

The last example works because you may omit the *on background* part from the syntax *foreground on background*, and just specify a foreground color. Epsilon interprets this as if you typed `on transparent`, and substitutes the background color specified for "text".

You can also specify that a particular color class is the same as a previously-defined color class, like this:

```
color_scheme "standard-gui" {
    color_class text = black on white;
    color_class tex_text = color_class text;
};
```

When, for the current scheme, there's no specific color information for a color class, Epsilon looks for a default color class specification, one that's not associated with any scheme:

```
color_class diff_added = black on yellow;
color_class c_string = cyan;
color_class c_charconst = color_class c_string;
```

The first definition above says that, in the absence of any color-scheme-specific setting for the `diff_added` color class, it should be displayed as black text on a yellow background. The second says that text in the `c_string` color class should be displayed using cyan text, on the default background for the scheme (that defined for the `text` color class). And the third says that text in the `c_charconst` color class should be displayed the same as text in the `c_string` color class for that scheme.

Internally, Epsilon stores all color class settings that occur outside any color scheme in a special color scheme, which is named "color-defaults". See page 493 for more on colors, including the other special-purpose flags that may appear in color class values.

9.13.3 Function Definitions

function-definition:

function-head block

```

function-head argument-decl-list block
ansi-function-head block
callable-function-head block
callable-function-head:
    typed-function-head
    command typed-function-head
    typed-function-head on binding-list
    command typed-function-head on binding-list
binding-list:
    keytable-name [ constant-expression ]
    keytable-name [ constant-expression ] , binding-list
keytable-name:
    identifier
typed-function-head:
    identifier ( )
    type-specifier identifier ( )
function-head:
    identifier ( argument-list )
    type-specifier identifier ( argument-list )
ansi-function-head:
    identifier ( ansi-argument-list )
    type-specifier identifier ( ansi-argument-list )
ansi-argument-list:
    type-specifier declarator
    type-specifier declarator , ansi-argument-list
argument-list:
    identifier
    identifier , argument-list
argument-decl-list:
    type-specifier declarator-list ;
    type-specifier declarator-list ; argument-decl-list

```

A function definition begins with a type specifier, the name of the function, and parentheses surrounding a comma-separated list of arguments. Any bindings may be given here using the `on` keyword, as described below. Declarations for the arguments then appear, and the body of the function follows. If the `command` keyword appears before the type specifier, the function is a command, and Epsilon will do completion on the function when it asks for the name of a command. A function may be a command only if it has no arguments.

You may omit the type specifier before the function name, in which case the function's type is `int`. You may also omit the declaration for any argument, in which case the argument will be an `int`. Note that unlike some languages such as Pascal, if there are no arguments, an empty pair of parentheses must still appear, both in the definition and where you call the function.

You may also define functions using ANSI C/C++ syntax, in which type information for function arguments appears with the argument names inside parentheses. These function headers have the same effect:

```

average(int count, short *values)    average(count, values)
                                     short *values;

```


When you call a function, arguments of type `char` or `short` are automatically changed to `ints`. A corresponding change happens to declarations of function arguments and return values. Additionally, function arguments declared as an array of some type are changed to be a pointer to the same type, just as array variables are changed to pointers to the start of the array when their names appear in expressions (see page 418). For example, these two function headers have the same effect.

```
short average(count, values)      average(count, values)
char count;                      short *values;
short values[ ];
```

The user can call any function which takes no arguments, or bind such a function to a key. Functions which are normally invoked in this way can be made commands with the `command` keyword, but this is not necessary. If you omit the `command` keyword, Epsilon will not perform command completion on the function's name. The `on` keyword can appear after the (empty) parentheses of a function's argument list, to provide bindings for the function. Each binding consists of a key table name, followed by a constant (the key number) in square brackets `[]`. There may be several bindings following the `on` keyword, separated by commas. You must have previously declared the key table name in the same file (or an `#included` file). The binding takes effect when you load the function.

Sometimes it is necessary to declare an identifier as a function, although the function is actually defined in a separately compiled source file. For example, you must declare a function before you use a pointer to that function. Also, the EEL compiler must know that a function returns a non-numeric type if its return value is used. Any declaration of an identifier with type *function returning . . .* is a function declaration. Function declarations may appear anywhere a local or global variable declaration is legal. So long as the identifier is not masked by a local variable of the same name, the declaration has effect until the end of the file.

Any function named `when_loading()` is automatically executed when you load the bytecode file it appears in into Epsilon. There may be any number of `when_loading()` functions defined in a file, and they execute in order, while the file is being loaded. Such functions are deleted as soon as they return. They may take no arguments.

9.14 Differences Between EEL And C

- Global variables may not be initialized with any expression involving pointers. This includes strings, which may only be used to directly initialize a declared array of characters. That is,

```
char example[ ] = "A string.;"
```

is legal, while

```
char *example = "A string.;"
```

is not.

- There are no static variables or functions. All local variables vanish when the function returns, and all global objects have names that separately compiled files can refer to.
- The C reserved word “extern” does not exist. In EEL, you may define variables multiple times with no problems, as long as they are declared to have the same type. The first definition read into Epsilon provides the initialization of the variable, and further initializations have no effect. However, if the variable is later declared with a different size, the size changes and the new initialization takes effect. To declare a function without defining it in a particular source file, see page 429.

- The C types “long”, “enum”, “void”, “float”, and “double” do not exist. Ints and shorts are always signed. Chars and bytes are always unsigned. There are no C bit fields. The C reserved words “long”, “float”, and “double” are not reserved in EEL.
- EEL provides the basic data type `spot`, and understands color class expressions and declarations using the `color_class` and `color_scheme` keywords.
- You may not cast between pointers and ints, except that function pointers may be cast to shorts, and vice versa. The constant zero may be cast to any pointer type. A pointer may be cast to a pointer of another type, with the exception of function pointers.
- You can use the reserved word `keytable` to declare empty key tables, as in

```
keytable reg_tab, cx_tab;
```

Local key tables are not permitted.

- The reserved word `command` is syntactically like a storage class. Use it to indicate that the function is normally called by the user, so `command` completion will work. The user can also call other functions (as long as they have no arguments) but the completion facility on `command` names ignores them.
- After the head of any function definition with no arguments, you can use the reserved word `on` to give a binding. It is followed by the name of a key table already declared, and an index (constant int expression) in square brackets. There may be more than one (separated by commas). For example,


```
command visit_file() on cx_tab[CTRL('V')]
```
- You can use the reserved word `buffer` as a storage class for global variables. It declares a variable to have a different value for each buffer, plus a default value. As you switch between buffers, a reference to a buffer-specific variable will refer to a different value.
- You can also use the reserved word `window` as a storage class for global variables. This declares the variable to have a different value for each window, plus a default value. As you switch between windows, a reference to a window-specific variable will refer to a different value.
- The reserved words `zeroed` and `user` do not exist in C. See page 424. The reserved word `volatile` does exist in ANSI C, but serves a different purpose in EEL. See page 411.
- The EEL statements `save_var`, `save_spot`, and `on_exit` do not exist in C. See page 414.
- In each compile, an include file with a certain name is only read once, even if there are several `#include` directives that request it.

9.15 Syntax Summary

program:

global-definition

global-definition program

global-definition:

function-definition

global-variable-definition

keytable-definition

typedef-definition

```

    color-class-definition
typedef-definition:
    typedef type-specifier declarator-list ;
color-class-definition:
    color_class color-class-list ;
    color_scheme color-scheme-list ;
color-class-list:
    color-class-item
    color-class-item , color-class-list
color-class-item:
    identifier
    identifier color_scheme string-constant = color-pair
    identifier { color-scheme-spec-list }
    identifier = color-pair
color-scheme-spec-list:
    color-scheme-spec
    color-scheme-spec color-scheme-spec-list
color-scheme-spec:
    color_scheme string-constant = color-pair ;
color-scheme-list:
    color-scheme-item
    color-scheme-item , color-scheme-list
color-scheme-item:
    string-constant
    string-constant color_class identifier = color-pair
    string-constant { color-class-spec-list }
color-class-spec-list:
    color-class-spec
    color-class-spec color-class-spec-list
color-class-spec:
    color_class identifier = color-pair ;
color-pair:
    color_class identifier
    constant-expression
    constant-expression on constant-expression
keytable-definition:
    keytable keytable-list ;
keytable-list:
    identifier
    identifier , keytable-list
global-variable-definition:
    type-specifier global-declarator-list ;
    global-modifier-list global-declarator-list ;
    global-modifier-list type-specifier global-declarator-list ;
global-modifier-list:

```

```

    global-modifier
    global-modifier global-modifier-list
global-modifier:
    buffer
    window
    zeroed
    user
    volatile
declarator-list:
    declarator
    declarator , declarator-list
declarator:
    identifier
    ( declarator )
    * declarator
    declarator [ constant-expression ]
    declarator [ ]
    declarator ( )
global-declarator-list:
    global-declarator
    global-declarator , global-declarator-list
global-declarator:
    declarator
    declarator = string-constant
    declarator = initializer
initializer:
    constant-expression
    string-constant
    { initializer-list }
    { initializer-list , }
initializer-list:
    initializer
    initializer , initializer-list
type-specifier:
    char
    short
    int
    struct struct-or-union-specifier
    union struct-or-union-specifier
    spot
    typedef-name
typedef-name:
    identifier
struct-or-union-specifier:
    struct-or-union-tag

```

```

    struct-or-union-tag { member-list }
    { member-list }
struct-or-union-tag:
    identifier
member-list:
    type-specifier declarator-list ;
    type-specifier declarator-list ; member-list
type-name:
    type-specifier abstract-declarator
abstract-declarator:
    empty
    ( abstract-declarator )
    * abstract-declarator
    abstract-declarator [ constant-expression ]
    abstract-declarator [ ]
    abstract-declarator ( )
    abstract-declarator ( ansi-argument-list )
function-definition:
    function-head block
    function-head argument-decl-list block
    ansi-function-head block
    callable-function-head block
callable-function-head:
    typed-function-head
    command typed-function-head
    typed-function-head on binding-list
    command typed-function-head on binding-list
binding-list:
    keytable-name [ constant-expression ]
    keytable-name [ constant-expression ] , binding-list
keytable-name:
    identifier
typed-function-head:
    identifier ( )
    type-specifier identifier ( )
function-head:
    identifier ( argument-list )
    type-specifier identifier ( argument-list )
ansi-function-head:
    identifier ( ansi-argument-list )
    type-specifier identifier ( ansi-argument-list )
ansi-argument-list:
    type-specifier declarator
    type-specifier declarator , ansi-argument-list
argument-list:

```

```

    identifier
    identifier , argument-list
argument-decl-list:
    type-specifier declarator-list ;
    type-specifier declarator-list ; argument-decl-list
block:
    { statement-list }
    { }
local-variable-definition:
    type-specifier local-declarator-list ;
local-declarator-list:
    local-declarator
    local-declarator , local-declarator-list
local-declarator:
    declarator
    declarator = expression
statement-list:
    statement
    statement statement-list
statement:
    expression ;
    if ( expression ) statement
    if ( expression ) statement else statement
    while ( expression ) statement
    do statement while ( expression );
    for ( opt-expression ; opt-expression ; opt-expression ) statement
    switch ( expression ) statement
    case constant-expression : statement
    default: statement
    break;
    continue;
    return;
    return expression ;
    save_var save-list ;
    save_spot save-list ;
    on_exit statement
    goto label ;
    label : statement
    ;
    local-variable-definition
    typedef-definition
    block
save-list:
    save-item
    save-item , save-list

```

```

save-item:
    identifier
    identifier = expression
    identifier modify-operator expression
    identifier ++
    identifier --
label:
    identifier
opt-expression:
    empty
    expression
expression:
    numeric-constant
    string-constant
    identifier
    identifier.default
    color_class identifier
    ( expression )
    ! expression
    * expression
    & expression
    - expression
    ~ expression
    sizeof expression
    sizeof( type-name )
    ( type-name ) expression
    ++ expression
    -- expression
    expression ++
    expression --
    expression + expression
    expression - expression
    expression * expression
    expression / expression
    expression % expression
    expression == expression
    expression != expression
    expression < expression
    expression > expression
    expression <= expression
    expression >= expression
    expression && expression
    expression || expression
    expression & expression
    expression | expression

```

expression ^ *expression*
expression << *expression*
expression >> *expression*
expression = *expression*
expression *modify-operator* *expression*
expression ? *expression* : *expression*
expression , *expression*
expression (
expression (*expression-list*)
expression [*expression*]
expression . *identifier*
expression -> *identifier*

modify-operator:

+=
 -=
 *=
 /=
 %=
 &=
 |=
 ^=
 <<=
 >>=

expression-list:

expression
expression , *expression-list*

constant-expression:

numeric-constant
 (*constant-expression*)
 ! *constant-expression*
 - *constant-expression*
 ~ *constant-expression*
 sizeof *constant-expression*
 sizeof (*type-name*)
constant-expression + *constant-expression*
constant-expression - *constant-expression*
constant-expression * *constant-expression*
constant-expression / *constant-expression*
constant-expression % *constant-expression*
constant-expression == *constant-expression*
constant-expression != *constant-expression*
constant-expression < *constant-expression*
constant-expression > *constant-expression*
constant-expression <= *constant-expression*
constant-expression >= *constant-expression*

constant-expression && constant-expression
constant-expression || constant-expression
constant-expression & constant-expression
constant-expression | constant-expression
constant-expression ^ constant-expression
constant-expression << constant-expression
constant-expression >> constant-expression
constant-expression ? constant-expression : constant-expression
constant-expression , constant-expression

Chapter 10

Primitives and EEL Subroutines



In this chapter, we describe all EEL primitives, as well as a few useful EEL subroutines. In Epsilon, the term “primitive” refers to a function or variable that is not written or defined in EEL, but rather built into Epsilon.

It’s possible for one command to call another command, as if it were a subroutine with no parameters, but commands are usually written to be called in the context of Epsilon’s main loop (see page 578), not from other commands. New Epsilon commands should generally call the functions defined in this chapter, not other commands, except in the case of a new command that simply adds some functionality to an existing single command.

Each section discusses items that pertain to a particular topic, and begins with EEL declarations for the items discussed in that section. If we implemented an item as an EEL subroutine, the declaration often includes a comment that identifies the EEL source file defining the item.

Some EEL primitives have optional parameters. For example, you can call the `get_tail()` primitive as either `get_tail(fname, 1)` or `get_tail(fname)`. Any missing parameter automatically takes a value of zero. In this manual, we indicate an optional parameter by showing a `?` before it.

When writing EEL extensions, an easy way to look up the documentation on the primitive or subroutine at point is to press F1 F (Enter). That displays the help in Info mode. Or you can use Shift-F1 to display it in a web browser.

10.1 Buffer Primitives

10.1.1 Changing Buffer Contents

```
insert(int ch)
user buffer int point;
```

An Epsilon *buffer* contains text that you can edit. Most of the primitives in this section act on, or refer to, one of the buffers designated as the *current buffer*.

The `insert()` primitive inserts a single character into the current buffer. Its argument says what character to insert. The buffer’s insertion point, or just point, refers to the particular position in each buffer where insertions occur.

The int variable named `point` stores this position. Its value denotes the number of characters from the beginning of the buffer to the spot at which insertions happen. For example, a value of zero for `point` means that insertions would occur at the beginning of the buffer. A value of one for `point` means that insertions would occur after the first character, etc.

To change the insertion point, you can assign a new value to `point`. For example, the statement

```
point = 3;
```

makes insertions occur after the third character in the buffer, assuming the buffer has at least 3 characters. If you set `point` to a value less than zero, `point` takes the value zero. Similarly, if you set `point` to a value greater than the size of the buffer, its value becomes the number of characters in the buffer.

When the current buffer changes, the value of the variable `point` automatically changes with it. We call variables with this behavior *buffer-specific* variables. See page 548.

```
int size()
```

The primitive function `size()` returns the number of characters in the current buffer. You cannot set the size directly: you can change the size of the buffer only by inserting or deleting characters. For this reason, we implemented `size()` as a function, not a variable like `point`.

The variable `point` refers not to a character position, but rather to a character boundary, a place between characters (or at the beginning or end of a buffer). The legal values for `point` range from zero to `size()`. We will refer to a value in this range, inclusive of the ends, as a *position*. A position is a place between characters in a buffer, or at the beginning of the buffer, or at the end. The value of a position is the number of characters before it in the buffer. In EEL, ints (integers) hold positions.

When Epsilon inserts a character, it goes before `point`, not after it. If Epsilon didn't work this way, inserting a, then b, then c would result in cba, not abc.

```
delete(int pos1, int pos2)
int delete_if_highlighted()
```

The `delete()` primitive deletes all characters between the two positions supplied as arguments to it. The order of the arguments doesn't matter.

The `delete()` primitive doesn't save deleted text in a kill buffer. The kill commands themselves manage the kill buffers, and use the `delete()` primitive to actually remove the text.

Commands that insert text often begin by calling the `delete_if_highlighted()` subroutine. If there's a highlighted region, this subroutine deletes it and returns 1. Otherwise (or if the `typing-deletes-highlight` variable has been set to zero), it returns 0.

```
replace(int pos, int ch)
int character(int pos)
int curchar()
```

The `replace()` primitive changes the character at position `pos` to `ch`. The parameter `pos` refers to the position before the character in question. Therefore, the value of `pos` can range from 0 to `size()-1`, inclusively.

The `character()` primitive returns the character after the position specified by its argument, `pos`. The `curchar()` returns the same value as `character(point)`. These two primitives return -1 when the position involved isn't valid, such as at the end of the buffer or before its start (when `pos` is less than zero). For example, `character(size())` returns -1, as does `curchar()` with `point` at the end of the buffer.

```
stuff(char *str)
int bprintf(char *format, ...)
int buffer_printf(char *name, char *format, ...)
int buf_printf(int bnum, char *format, ...)
buf_stuff(int bnum, char *s, int len)
```

The `stuff()` function inserts an entire string into the current buffer.

The `bprintf()` function also inserts a string, but it takes a format string plus other arguments and builds the string to insert using the rules on page 480. The `buffer_printf()` function operates similarly, except that it takes the name of the buffer into which to insert the string. It creates the buffer if necessary. Similarly, `buf_printf()` takes a buffer number, and inserts the formatted string into that buffer. All of the primitives described in this paragraph return the number of characters they inserted into the buffer.

The `buf_stuff()` primitive includes a length parameter, so it can handle text that may contain null characters. It inserts the `len` characters at `s` into the buffer `bnum`.

10.1.2 Moving Text Between Buffers

```
xfer(char *buf, int from, int to)
buf_xfer(int bnum, int from, int to) /* buffer.e */
raw_xfer(int bnum, int from, int to)
buf_xfer_colors(int bnum, int from, int to)
grab_buffer(int bnum) /* buffer.e */
```

The `xfer()` subroutine transfers characters from one buffer to another. It copies the characters between `from` and `to` in the current buffer and inserts them at `point` in the named buffer. It positions the mark in the named buffer just before the inserted characters, and positions its `point` right after the insertion. The current buffer doesn't change. The `buf_xfer()` subroutine works similarly, but accepts a buffer number instead of a name. Both use the `raw_xfer()` primitive to transfer the text.

The `buf_xfer_colors()` subroutine is like `buf_xfer()`, but copies any colors set by `set_character_color()` as well.

The `grab_buffer()` subroutine copies text in the other direction. It inserts the text of buffer number `bnum` into the current buffer before `point`, setting the mark before the inserted text.

10.1.3 Getting Text from a Buffer

```
grab(int pos1, int pos2, char *to)
grab_expanding(int pos1, int pos2, char **toptr, int minlen)
buf_grab_bytes(int buf, int from, int to, char *dest)
```

The `grab()` primitive copies characters from the buffer to a string. It takes the range of characters to copy, and a character pointer indicating where to copy them. The buffer doesn't change. The positions may be in either order. The resulting string will be null-terminated.

The `grab_expanding()` subroutine is similar, but works with a dynamically allocated character pointer, not a fixed-length character array. Pass a pointer to a `char *` variable, and the subroutine will resize it as needed to hold the result. The `char *` variable may hold NULL initially. The `minlen` parameter provides a minimum allocation length for the result.

The `buf_grab_bytes()` subroutine copies characters in the specified range in the buffer `buf` into the character array `dest`, in the same fashion as `grab()`. Despite its name, it operates on 16-bit characters, not 8-bit bytes.

```
grab_full_line(int bnum, char *str) /* buffer.e */
grab_line(int bnum, char *str) /* buffer.e */
```

The `grab_full_line()` subroutine copies the entire current line of buffer number `bnum` into the character array `str`. It doesn't change `point`. The `grab_line()` subroutine copies the remainder of `bnum`'s current line to `str`, and moves to the start of the next line. Neither function copies the `<Newline>` at the end of the line, and each returns the number of characters copied.

```
grab_line_offset(int b, char *s, int offset, int wrap)
get_line_from_buffer(int buf, int line, char *res)
move_line_to_buffer(int dest)
```

The `grab_line_offset()` subroutine copies a line from buffer `b` into `s`, discarding any trailing newline and returning the string's length. The `offset` parameter specifies which line: 0 means the buffer's current line, 1 means the next, and so forth. The subroutine moves to end of the appropriate line. If `offset` is negative, the subroutine moves to a previous line and copies it, remaining at its start. The `wrap` parameter, if nonzero, makes the subroutine wrap around to the opposite end of the buffer if it hits an end when counting lines; if zero, a too-high offset returns zero and empties `s`.

The `get_line_from_buffer()` subroutine copies the `line`'th line in buffer `buf` to `res`.

The `move_line_to_buffer()` subroutine copies the current line to the buffer `dest`, then deletes it from the current buffer.

```
int grab_numbers(int bnum, int *nums)    /* buffer.e */
int break_into_numbers(char *s, int *nums)
```

The `grab_numbers()` subroutine uses `grab_line()` to retrieve a line from buffer `bnum`. Then it breaks the line into words (separated by spaces and tabs), and tries to interpret each word as a number by calling the `numtoi()` subroutine. It puts the resulting numbers in the array `nums`. The function returns the number of words on the line.

The `break_into_numbers()` subroutine is similar, but retrieves the numbers from a string `s`. It returns the count of numbers it found, putting their values into the `nums` array.

```
int grab_string(int bnum, char *s, char *endmark) /* buffer.e */
int grab_string_expanding(int bnum, char **s,
                          char *endmark, int minlen)
```

The `grab_string()` subroutine copies from buffer `bnum` into `s`. It copies from the buffer's current position to the beginning of the next occurrence of the text `endmark`, and leaves the buffer's point after that text. It returns 1, unless it couldn't find the `endmark` text. In that case, it moves to the end of the buffer, sets `s` to the empty string, and returns 0.

The `grab_string_expanding()` subroutine is similar, but works with dynamically allocated character pointers, not fixed-length character arrays. Pass a pointer to a `char *` variable, and the subroutine will resize it as needed to hold the result. The `char *` variable may hold NULL initially. The `minlen` parameter provides a minimum allocation length for the result.

10.1.4 Spots

```
spot alloc_spot(?int left_ins)
free_spot(spot sp)
int spot_to_buffer(spot sp)
```

A place in the buffer is usually recorded and saved for later use as a count of the characters before that place: this is a position, as described on page 440. Sometimes it is important for the stored location to remain between the same pair of characters even if many changes are made to other parts of the buffer (affecting the number of characters before the saved location).

Epsilon provides a type of variable called a *spot* for this situation. The declaration

```
spot sp;
```

says that `sp` can refer to a spot. It doesn't create a new spot itself, though.

The `alloc_spot()` primitive creates a new spot and returns it, and the `free_spot()` primitive takes a spot and discards it. The spot that `alloc_spot()` returns is initially set to `point`, and is associated with the current buffer. Deleting a buffer frees all spots associated with it. If you try to free a spot whose buffer has already been deleted, Epsilon will ignore the request, and will not signal an error.

The `spot_to_buffer()` primitive takes a spot and returns the buffer number it was created for, or `-1` if the buffer no longer exists, or `-2` if the buffer exists, but that particular spot has since been deleted.

If the `left_ins` parameter to `alloc_spot()` is nonzero, a left-inserting spot is created. If the `left_ins` parameter is 0, or is omitted, a right-inserting spot is created. The only difference between the two types of spots is what they do when characters are inserted right where the spot is. A left-inserting spot stays after such inserted characters, while a right-inserting spot stays before them. For example, imagine an empty buffer, with all spots at 0. After five characters are inserted, any left-inserting spots will be at the end of the buffer, while right-inserting spots will remain at the beginning.

A spot as returned by `alloc_spot()` behaves a little like a pointer to an int, in that you must dereference it by writing `*sp` to obtain the position it currently refers to. For example:

```
fill_all()      /* fill paragraphs, leave point alone */
{
    spot oldpos = alloc_spot(), oldmark = alloc_spot();

    *oldpos = point;
    *oldmark = mark;      /* save old values */
    point = 0;           /* make region be whole buffer */
    mark = size();
    fill_region();       /* fill paragraphs in region */
    mark = *oldmark;     /* restore values */
    point = *oldpos;
    free_spot(oldmark);  /* free saving places */
    free_spot(oldpos);
}
```

A simpler way to write the above subroutine uses EEL's `save_spot` keyword. The `save_spot` keyword takes care of allocating spots, saving the original values, and restoring those values when the subroutine exits. See page 414 for more on `save_spot`.

```
fill_all()      /* fill paragraphs, leave point alone */
{
    /* uses save_spot */
    save_spot point = 0; /* make region be whole buffer */
    save_spot mark = size();
    fill_region();      /* fill paragraphs in region */
}
```

Like a pointer, a spot variable can contain zero, and `alloc_spot()` is guaranteed never to return this value. Epsilon signals an error if you try to dereference a spot which has been freed, or whose buffer no longer exists.

```
buffer spot point_spot;
buffer spot mark_spot;
#define point    *point_spot
```

```
#define mark      *mark_spot
/* These variables are actually defined
   differently. See below. */
```

Each new buffer begins with two spots, `point_spot` and `mark_spot`, set to the beginning of the buffer. `Point_spot` is a left-inserting spot, while `mark_spot` is a right-inserting spot. These spots are created automatically with each new buffer, and you cannot free them. You can think of the built-in variables `point` and `mark` as simply macros that yield `*point_spot` and `*mark_spot`, respectively. That's why you don't need to put a `*` before each reference to `point`.

```
user buffer int point;      /* True definitions */
user buffer int mark;
spot get_spot(int which)
#define point_spot      get_spot(0)
#define mark_spot      get_spot(1)
```

Actually, while `point` and `mark` could be defined as macros, as above, they're not. Epsilon recognizes them as built-in primitives for speed. On the other hand, `point_spot` and `mark_spot` actually are macros! They use the `get_spot()` primitive, which has no function other than to return these two values.

```
do_set_mark(int val)
```

The `do_set_mark()` subroutine sets the current buffer's mark to the specified value. It also records the current virtual column (which, typically, should match the mark). The rectangle commands retrieve this, so that in virtual mode you can copy rectangles that end in virtual space.

```
set_spot(spot *s, int pos)
```

The `set_spot()` subroutine sets a spot so it refers to the position `pos` in the current buffer. Pass it the address of a spot variable. If the spot is zero or refers to a different buffer, the subroutine will create a new right-inserting spot in the current buffer, freeing the old spot.

10.1.5 Narrowing

```
user buffer int narrow_start;
user buffer int narrow_end;
int narrow_position(int p)      /* buffer.e */
```

Epsilon provides two primitive variables, `narrow_start` and `narrow_end`, that restrict access to the current buffer. The commands `narrow-to-region` and `widen-buffer`, described on page 176, use these variables. Epsilon ignores the first `narrow_start` characters and the last `narrow_end` characters of the buffer. Usually, these variables have a value of zero, so no such restriction takes place. Characters outside of the narrowed region will not appear on the screen, and will remain outside the control of normal Epsilon commands.

If you try to set a primitive variable such as `point` to a position outside of the narrowed area, Epsilon will change the value to one inside the narrowed area. For example, suppose the buffer contains one hundred characters, with the first and last ten characters excluded, so only eighty appear on the screen. In this case, `size()` will return one hundred, and `narrow_start` and `narrow_end` will each have a value of ten. The statement `point = 3;` will give `point` a value of ten (the closest legal value), while the statement `point =`

10000; will give point the value ninety. Epsilon adjusts the parameters of primitive functions in the same way. Suppose, in the example above, you try to delete all the characters in the buffer, using the `delete()` primitive. Epsilon would take the statement `delete(0, size());` and effectively change it to `delete(10, 90);` to delete only the characters inside the narrowed area.

The `narrow_position()` subroutine returns its argument `p`, adjusted so that it's inside the narrowed buffer boundaries.

Writing the buffer to a file ignores narrowing. Reading a file into the buffer lifts any narrowing in effect by setting `narrow_start` and `narrow_end` to zero.

10.1.6 Undo

```
int undo_op(int is_undo)
undo_mainloop()
undo_redisplay()
user buffer int undo_size;
```

With a nonzero argument, the `undo_op()` primitive undoes one basic operation like the `undo` command, described on page 105. With an argument of zero, it acts like `redo`. It returns a bit pattern describing what types of operations were undone or redone. The bit codes are defined in `codes.h`. `UNDO_INSERT` means that originally an insertion occurred, and it was either undone or redone. The `UNDO_DELETE` and `UNDO_REPLACE` codes are similar.

Epsilon groups individual buffer changes into groups, and undoes one group at a time. While saving changes for undoing, Epsilon begins a new group when it redisplays buffers or when it begins a new command in the main loop. The `UNDO_REDISP` code indicates the former happened, and `UNDO_MAINLOOP` the latter. `UNDO_MOVE` indicates movement is being undone, and `UNDO_END` is used when Epsilon could only undo part of a command. If `undo_op()` returns zero, the buffer was not collecting undo information (see below).

Epsilon automatically starts a new undo group each time it does normal redisplay or passes through its main loop, by calling either the `undo_redisplay()` or `undo_mainloop()` primitives, respectively. You can call either of these primitives yourself to make Epsilon start a new undo group.

In addition to starting a new group, the `undo_mainloop()` primitive also makes the current buffer start to collect undo information. When you first create a buffer, Epsilon doesn't keep undo information for it, so that "system" buffers don't have this unnecessary overhead. Each time it passes through the main loop, Epsilon calls `undo_mainloop()`, and this makes the current buffer start collecting undo information, if it isn't already, and if the buffer-specific variable `undo_size` is nonzero.

```
int undo_count(int is_undo)
```

The `undo_count()` primitive takes a parameter that specifies whether undoing or redoing is involved, like `undo_op()`. The primitive returns a value indicating how much undoing or redoing information is saved. The number doesn't correspond to a particular number of commands, but to their complexity.

```
user buffer int undo_flag;
```

In addition to buffer changes and movements, Epsilon can record other information in its list of undoable operations. Each time you set the `undo_flag` variable, Epsilon inserts a "flag" in its undo list with the particular value you specify. When Epsilon is undoing or redoing and encounters a flag, it immediately ends the current group of undo operations and returns a code with the `UNDO_FLAG` bit on. It puts the value of the flag it encountered in the `undo_flag` variable. The `yank-pop` command uses flags 1 and 2 for undoing the previous `yank`.

10.1.7 Searching Primitives

```

user int matchstart;
user int matchend;
int search(int dir, char *str)
user short abort_searching;
#define ABORT_JUMP      -1
#define ABORT_ERROR     -2

```

The search primitives each look for the first occurrence of some text in a particular direction from point. Use 1 to specify forward, -1 to specify backward. They move point to the far end of the match, and set the matchstart and matchend variables to the near and far ends of the match, respectively. For example, if the buffer contains “abcd” and you search backward from the end for “bc”, point and matchend will be 1 (between the ‘a’ and the ‘b’) and matchstart will be 3. If the search text does not appear in the buffer, point goes to the appropriate end of the buffer. These primitives return 1 if they find the text and 0 if not.

The most basic searching function is the `search()` primitive. It takes a direction and a string, and searches for the string. It returns 1 if it finds the text, or 0 if it does not.

If the user presses the abort key during searching, Epsilon’s behavior depends upon the value of the `abort_searching` variable. If it’s `ABORT_IGNORE` (0), the key is ignored and the search continues. If it’s `ABORT_JUMP` (the default), Epsilon aborts the search and jumps by calling the `check_abort()` primitive. If it’s `ABORT_ERROR`, Epsilon aborts the search and returns the value `ABORT_ERROR`. The `search()`, `re_search()`, `re_match()`, and `buffer_sort()` primitives all use the `abort_searching` variable to control aborting.

```

user buffer short case_fold;

```

If the `case-fold` buffer-specific variable is nonzero, characters that match except for case count as a match. Otherwise, only exact matches (including case) count. To alter folding rules, see page 538.

Regular Expression Searching

```

int re_search(int flags, char *pat)
int re_compile(int flags, char *pat)
int re_match()
#define RE_FORWARD      0
#define RE_REVERSE     2
#define RE_FIRST_END   4
#define RE_SHORTEST    8
#define RE_IGNORE_COLOR 16

```

Several searching primitives deal with a powerful kind of pattern known as a *regular expression*. Regular expressions allow you to search for complex patterns. Regular expressions are strings formed according to the rules on page 66.

The `re_search()` primitive searches the buffer for one of these patterns. It operates like the `search()` primitive, taking a direction and pattern and returning 1 if it finds the pattern. It moves to the far end of the pattern from the starting point, and sets matchstart to the near end. If it doesn’t find the pattern, or if the pattern is illegal, it returns 0. In the latter case point doesn’t move, in the former point moves to the end (or beginning) of the buffer.

When you specify a direction using 1 or -1, Epsilon selects the first-beginning, longest match, unless the search string overrides this. However, instead of providing a direction (1 or -1) as the first parameter to `re_search()` or `re_compile()`, you can provide a set of flags. These let you specify finding the shortest possible match, for example, without altering the search string.

The `RE_FORWARD` flag searches forward, while the `RE_REVERSE` flag searches backward. (If you don't include either, Epsilon searches forward.) The `RE_FIRST_END` flag says to find a match that ends first, rather than one that begins first. The `RE_SHORTEST` flag says to find the shortest possible match, rather than the longest. However, if the search string contains sequences that specify first-ending, first-beginning, shortest, or longest matches, those sequences override any flags.

A pattern may include color class assertions, as described on page 74. The `RE_IGNORE_COLOR` flag makes Epsilon ignore such assertions. The `do_color_searching()` subroutine uses this; if your search might include such assertions, calling that subroutine instead of these primitives will take care of ensuring that the buffer's syntax highlighting is up to date.

The `re_compile()` primitive checks a pattern for legality. It takes the same arguments as `re_search()` and returns 1 if the pattern is illegal, otherwise 0. The `re_match()` primitive tells if the last-compiled pattern matches at this location in the buffer, returning the far end of the match if it does, or -1 if it does not.

```
int parse_string(int flags, char *pat, ?char *dest)
int matches_at(int pos, int dir, char *pat)
int matches_at_length(int pos, int dir, char *pat)
int matches_in(int start, int end, char *pat)
```

The `parse_string()` primitive looks for a match starting at point, using the same rules as `re_match()`. It takes a direction (or flags) and a pattern like `re_compile()`, and a character pointer. It looks for a match of the pattern beginning at point, and returns the length of such a match, or zero if there was no match.

The third argument `dest` may be a null pointer, or may be omitted entirely. But if it's a pointer to a character array, `parse_string()` copies the characters of the match there, and moves point past them. If the pattern does not match, `dest` isn't modified.

The `matches_at()` subroutine accepts a regular expression `pat` and returns nonzero if the given pattern matches at a particular position in the buffer in the given direction. The `matches_at_length()` subroutine is similar, but it returns the length of the match, or zero if there was no match.

The `matches_in()` subroutine accepts a regular expression `pat` and searches for the pattern in the specified buffer range, returning nonzero to indicate it matches. Neither `matches_at()` nor `matches_in()` move point.

```
int find_group(int n, int open)
```

The `find_group()` primitive tells where in the buffer certain parts of the last pattern matched. It counts opening parentheses used for grouping in the last pattern, numbered from 1, and returns the position it was at when it reached a certain parenthesis. If `open` is nonzero, it returns the position of the `n`'th left parenthesis, otherwise it returns the position of its matching right parenthesis. If `n` is zero, it returns information on the whole pattern. If `n` is too large, or negative, the primitive aborts with an error message. Parentheses that use the syntax `(?:)` don't count.

Searching Subroutines

```
int do_searching(int flags, char *str) /* search.e */
int do_color_searching(int flags, char *str) /* search.e */
```

The `do_searching()` subroutine defined in `search.e` is handy when you want to use a variable to determine the type of search. A `flags` value of 0 means perform a plain forward search. The flags `REVERSE`, `REGEX`, and `WORD` specify a reverse search, a regular expression search, or a word search, respectively. The subroutine normally performs case-folding if the buffer's `case_fold` variable is non-zero; pass `MODFOLD` to force Epsilon to search without case-folding, or pass `MODFOLD` and `FOLD` to force Epsilon to case-fold. The above flags may be combined in any combination.

The `do_searching()` subroutine returns 1 on a successful search, or 0 if the search text was not found. It can also return `DSABORT` if the user aborted the search (see the `abort_searching` variable) or `DSBAD` if the (regular expression) search pattern was invalid. If the search was successful, Epsilon moves to just after the found text (or just before, for reverse searches); in all other cases point doesn't change.

The `do_color_searching()` subroutine defined in `search.e` takes parameters and returns values just like `do_searching()`, but it handles regular expressions that use assertions like `<c:perl-comment>` to match based on the colors applied via syntax highlighting. If you use such syntax in a primitive like `re_search()`, Epsilon will search based on the syntax highlighting currently applied to the buffer. Because Epsilon computes syntax highlighting only as needed during screen display, as well as in the background, the buffer's syntax highlighting may not be up to date. This subroutine ensures that the buffer's syntax highlighting is up to date as it finds matches, by reparsing and recoloring the buffer whenever it has to.

```
int word_search(int dir, char *str)
int narrowed_search(int flags, char *str, int limit)
```

If `do_searching()` needs to search in word mode, it calls the `word_search()` subroutine. This function searches for `str`, rejecting matches unless they are preceded and followed by non-word characters. More precisely, it converts the text into a regular expression pattern, constructed so that each space in the original pattern matches any sequence of whitespace characters, and each word in the pattern only matches whole words.

The `narrowed_search()` subroutine is like `do_searching()`, but takes a parameter `limit` to limit the search. Epsilon will only search a region of the buffer within `limit` characters of its starting point. For example, if point is at 30000 and you call `narrowed_search()` and specify a reverse search with a limit of 1000, the match must occur between positions 29000 and 30000. If no such match is found, point will be set to 29000 and the function will return 0.

```
string_replace(char *str, char *with, int flags)
show_replace(char *str, char *with, int flags)
```

The `string_replace()` subroutine allows you to do string replacements from within a function. It accepts flags from the same list as `do_searching()`. Provide the `INCR` flag if you want the subroutine to display the number of matches it found, and the number that were replaced. Provide the `QUERY` flag to ask the user to confirm each replacement. This subroutine sets the variables `replace-num-found` and `replace-num-changed` to indicate the total number of replacements it found, and the number the user elected to change.

If you want to display what will be replaced without replacing anything, call the `show_replace()` subroutine. It takes the same parameters as `string_replace()`, and displays a message in the echo area. All Epsilon's replacing commands call this subroutine to display their messages.

```
int simple_re_replace(int flags, char *str, char *repl)
```

The `simple_re_replace()` subroutine performs a regular expression replacement on the current buffer. It searches through the buffer, starting from the top (the bottom, for reverse searches), and passing flags and `str` directly to the `re_search()` primitive. It deletes each match and inserts the string `repl` instead, returning the number of replacements it did. The replacement text is inserted literally, with no interpolation. If you want to use `#1` in your replacement text or other more involved things, call `string_replace()` instead. The subroutine preserves point using a spot; if the text containing point is replaced, point will go after the replacement.

```
int search_read(char *str, char *prmt, int flags)
int default_fold(int flags)
int get_search_string(char *pr, int flags)
char *default_search_string(int flags)
char **default_replace_string(int flags)
```

To ask the user for a search string, use the `search_read()` subroutine. Its parameter `str` provides an initial search string, and it returns a set of flags which you can pass to `do_searching()`. It takes an initial set of flags, which you can use to start the user in one of the searching modes. Call `default_fold()` with any flags before calling `search_read()`. It will turn on any needed flags relating to case-folding, based on the value of the `case_fold` variable, and return a modified set of flags.

The function leaves the string in either the `_default_search` or the `_default_regex_search` variable, depending upon the searching flags it returns. You can call the `default_search_string()` subroutine with that set of searching flags and it will return a pointer to the appropriate one of these. Depending on what the user types, the `search_read()` subroutine may perform searching itself, in addition to returning the search string.

The similar `default_replace_string()` subroutine returns a pointer to the address of the current replacement string. Dereference it to access the replacement string.

The `get_search_string()` subroutine asks the user for a string to search for by calling `search_read()`.

```
buffer int (*search_continuation)();
int sample_search_continuation(int code, int flags, char *str)
```

In some modes a buffer may contain a single “record” out of many. Records may be swapped by changing the narrowing on the buffer (as in Info mode), while in other modes the contents of the buffer may be completely replaced with text from a different record.

A mode may wish to let users search from one record to the next, when no more matches can be found in the current record. (This capability relates to searching by the user, with the `search_read()` subroutine, not the primitive searching functions.)

A mode may set the buffer-specific `search_continuation` function pointer to a search-continuation function if it wants this behavior. If it’s nonzero, the searching functions will call this function to advance to a different record, or to remember or return to a particular record.

Epsilon assumes that the set of possible records have an implicit order to them, forming a list. And it assumes that a record id, referring to a specific record, may be stored in a character array of length `FNAMELEN`.

The code parameter indicates the desired operation. If `SCON_RECORD`, the search-continuation function must write a record id for the current record into the array `str`. If `SCON_RESTORE`, it must return to the

record identified by the previously-saved `id` `str`. These operations should return zero. If `SCON_COMPARE`, it must compare the current record with the `id` saved in `str` (according to the record order), returning -1, 0, or 1 depending on whether the current record is before, equal to, or after the saved record, respectively.

Any other code means to move to the next or previous record, according to whether the `flags` parameter contains the `REVERSE` bit, and position to its start (or, for reverse searching, end). In this case, code becomes a count, starting from 1, that indicates the number of record positionings done since the last user keypress (for use in displaying progress messages). It should return 1 on success, or 0 if there were no more records (and should remain at the original record in that case).

A search-continuation function may wish to pre-screen records, and skip over those that do not contain the search string (but is not required to do so). If it chooses to do this, it can use `flags` and `str` to call the `do_searching()` subroutine; these specify the search being performed.

```
int col_search(char *str, int col)    /* search.e */
int column_color_searching(int flags, char *pat, int startcol, int endcol)
```

The `col_search()` subroutine defined in `search.e` attempts to go to the beginning of the next line containing a certain string starting in a certain column. It returns 1 if the search is successful, 0 otherwise.

The `column_color_searching()` subroutine defined in `search.e` ignores matches unless they start and end in the specified columns. Either `startcol` or `endcol` may be -1, and that column restriction won't apply. It takes `flags` and `pat` parameters like `do_color_searching()`, so it can search for regular expressions, including those that use syntax highlighting colors, restrict matches to whole words, search in reverse, and so forth. See `do_searching()` at the start of this section for details on its flag parameter.

```
int line_search(int dir, char *s)    /* grep.e */
int prox_line_search(char *s)       /* tags.e */
int count_matches_in_buffer(int buf, char *text)
```

The `line_search()` subroutine searches in direction `dir` for a line containing only the text `s`. It returns 1 if found, otherwise 0.

The `prox_line_search()` subroutine searches in the buffer for lines containing exactly the text `s`. It goes to the start of the closest such line to point, and returns 1. If there is no matching line, it returns 0.

The `count_matches_in_buffer()` subroutine returns the number of instances of the given text in the specified buffer.

```
do_drop_matching_lines(int flags, char *pat, int drop)
```

The `do_drop_matching_lines()` subroutine deletes all lines after point in the current buffer but those that contain the specified search pattern. The search flags say how to interpret the pattern. If `drop` is nonzero, the subroutine deletes lines that contain the pattern; if `drop` is zero it deletes all lines except those that contain the pattern. Temporarily set the `show-status` variable to zero to keep it from displaying a line count summary.

```
replace_in_readonly_hook(int old_readonly)
replace_in_existing_hook(int old_readonly)
```

The `file-query-replace` command calls some hook functions as it goes through its list of buffers or files. Just before it makes its first change in each buffer (or asks the user whether to make the change, if it's still in query mode), it calls either the `replace_in_existing_hook()` subroutine (if the buffer or file was

already loaded before running the command) or the `replace_in_readonly_hook()` (if file-query-replace had to read the file itself). The file-query-replace command temporarily zeroes the `readonly-warning` variable; it passes the original value of this variable as a parameter to each hook.

The default version of `replace_in_existing_hook()` does nothing. The default version of `replace_in_readonly_hook()` warns about the file being read-only by calling `do_readonly_warning()`.

10.1.8 Moving by Lines

```
int nl_forward()
int nl_reverse()
int move_by_lines(int cnt)
```

The `nl_forward()` and `nl_reverse()` primitives quickly search for newline characters in the direction you specify. The `nl_forward()` primitive is the same as `search(1, "\n")`, while `nl_reverse()` is the same as `search(-1, "\n")`, where `\n` means the newline character (see page 402). These primitives do not set `matchstart` or `matchend`, but otherwise work the same as the previous searching primitives, returning 1 if they find a newline and 0 if they don't.

The `move_by_lines()` primitive moves forward over `cnt` lines, like calling `nl_forward()` that many times. If `cnt` is negative, it moves backward, like calling `nl_reverse()`. It returns 0, unless it hit the end of the buffer (or narrowing) before moving the full amount; in that case it returns the number of lines still to go when it stopped. For example, if there are two newlines in the buffer before point, calling `move_by_lines(-10)` moves to the start of the buffer and returns -8.

```
to_begin_line()      /* eel.h macro */
to_end_line()        /* eel.h macro */
int give_begin_line() /* basic.e */
int give_end_line()   /* basic.e */
```

The `eel.h` file defines textual macros named `to_begin_line()` and `to_end_line()` that make it easy to go to the beginning or end of the current line. They simply search in the appropriate direction for a newline character and back up over it if the search succeeds.

The `give_begin_line()` subroutine returns the buffer position of the beginning of the current line, and the `give_end_line()` subroutine returns the position of its end. Neither moves point.

```
go_line(int num)      /* basic.e */
buf_go_line(int buf, int num)
int lines_between(int from, int to, ?int abort_ok)
int count_lines_in_buf(int buf, int abortok)
int buf_position_to_line_number(int buf, int pos)
int all_blanks(int from, int to) /* indent.e */
```

The EEL subroutine `go_line()` defined in `basic.e` uses the `move_by_lines()` primitive to go to a certain line in the buffer, counting from 1. `go_line(2)`, for example, goes to the beginning of the second line in the buffer. The similar `buf_go_line()` subroutine does the same in the specified buffer.

The `lines_between()` primitive returns the number of newline characters in the part of the buffer between `from` and `to`. If `abort_ok` is nonzero, the user can abort from this primitive, otherwise Epsilon ignores the abort key.

The `buf_position_to_line_number()` subroutine returns the line number, counting from 1, of a particular position in the specified buffer.

The `count_lines_in_buf()` subroutine returns the number of newline characters in the buffer `buf`. If `abortok` is nonzero and the user press the abort key, the subroutine uses the `check_abort()` primitive to abort.

The `all_blanks()` subroutine returns 1 if the characters between `from` and `to` are all whitespace characters (space, tab, or newline), 0 otherwise.

10.1.9 Other Movement Functions

```
int move_level(int dir, char *findch,
               char *otherch, int show, int stop_on_key)
buffer int (*mode_move_level)();
int c_move_level(int dir, int stop_on_key)
int html_move_level(int dir, int stop_on_key)
int default_move_level(int dir, char *findch,
                       char *otherch)
```

Several subroutines move through text counting and matching various sorts of delimiters. The `move_level()` subroutine takes a direction `dir` which may be 1 or -1, and two sets of delimiters. The routine searches for any one of the characters in `findch`. Upon finding one, it continues searching in the same direction for the character in the same position in `otherch`, skipping over matched pairs of these characters in its search.

For example, if `findch` was ">]" and `dir` was -1, `move_level()` would search backwards for one of these three characters. If it found a '[' first, it would then select the third character of `otherch`, which might be a '('. It would then continue searching for a '('. But if it found additional '[' characters before reaching that '(', it would need to find additional '(' characters before stopping.

The subroutine returns 1 to indicate that it found a match, and leaves `point` on the far side of the match (like commands such as `forward-level`). If no match can be found, the subroutine returns 0. Additionally, if its parameter `show` is nonzero, it displays an "Unmatched delimiter" message. When no characters in `findch` can be found in the specified direction, it sets `point` to the far end of the buffer and returns 1. If `stop_on_key` is nonzero, the subroutine will occasionally check for user key presses, and abort its search if the user has pressed a key. It returns -2 in this case and doesn't change `point`.

Certain modes define a replacement level matcher that understands more of the syntax of that mode's language. They do this by setting the buffer-specific function pointer variable `mode_move_level` to a function such as `c_move_level()`. The `move_level()` subroutine will call this function instead of doing its normal processing when this variable is nonzero in the current buffer.

Any such function will receive only `dir` and `stop_on_key` parameters. (It should already know which delimiters are significant in its language.) It should return the buffer position it reached (but not actually move there), if it found a pair of matched delimiters, or if it reached one end of the buffer without finding any suitable delimiters. It should return -1 if it detected an unmatched delimiter, or -2 if a keypress made it abort.

The `default_move_level()` function is what `move_level()` calls when no mode-specific function is available. It takes parameters like `move_level()`, and returns -1 or a buffer position like `c_move_level()`. A mode-specific function may wish to call this function, specifying a set of delimiters suitable for that language. The `html_move_level()` subroutine, for example, does just that.

```
int give_position(int (*cmd)())
```


The `give_position()` subroutine runs the subroutine `cmd`, which (typically) moves to a new position in the buffer. The `give_position()` subroutine returns this new position, but restores point to its original value. For example, `give_position(forward_word)` returns the buffer position of the end of the current word. EEL requires that `cmd` be declared before you call it, via a line like `int cmd();`, unless it's defined in the same file, before the `give_position()` call.

10.1.10 Sorting Primitives

```
buffer_sort(char *newbuf, ?int col, int rev)
do_buffer_sort(char *newbuf, int col, int rev)
sort_another(char *buf, int col, int rev)
do_sort_region(int from, int to, int col, int rev)
char show_status;
```

The EEL primitive `buffer_sort()` sorts the lines of the current buffer alphabetically. It does not modify the buffer, but rather inserts a sorted copy into the named buffer (which must be different). It performs each comparison starting at column `col`, which is optional and defaults to 0 (the first column). If the `case_fold` variable is nonzero, sorting ignores the case of letters. It sorts lines in reverse order if the optional `rev` parameter is 1, not 0. (To alter folding rules, see page 538.)

If the variable `show_status` is nonzero, Epsilon will display progress messages as the sort progresses. Otherwise, no status messages appear.

The `do_buffer_sort()` subroutine is similar, but respects the `sort-case-fold` variable, not `case-fold` like `buffer_sort()`.

The `sort_another()` subroutine takes the name of a buffer and sorts it in place. The parameter `col` specifies the column to sort on, and `rev`, if nonzero, requests a reverse sort.

The `do_sort_region()` subroutine sorts a portion of the current buffer in place. The `from` and `to` parameters specify the region to sort. The `col` parameter specifies the column to sort on, and the `rev` parameter, if nonzero, requests a reverse sort.

If the user presses the abort key during sorting, Epsilon's behavior depends upon the value of the `abort_searching` variable. If 0, the key is ignored and the sort will run to completion. If `ABORT_JUMP`, Epsilon aborts the sort and jumps by calling the `check_abort()` primitive. If `ABORT_ERROR`, Epsilon aborts the sort and returns `ABORT_ERROR`. Whenever Epsilon aborts a sort, nothing gets inserted in the `newbuf` buffer. (For the subroutines that sort in place, the buffer is not changed.) Except when aborted, the `buffer_sort()` primitive and all the sorting subroutines described above return 0.

10.1.11 Other Formatting Functions

```
int unicode_character_name(char *name, int ch, ?int trail_ch)
#define IS_LEADING_SURROGATE(c)      (((c)&0xffffc00)==0xd800)
#define IS_TRAILING_SURROGATE(c)    (((c)&0xffffc00)==0xdc00)
```

The `unicode_character_name()` primitive converts between the name of a Unicode character and its numeric code. It may be used in several ways.

When `ch` is the code for a Unicode character and `trail_ch` is zero, the primitive writes the character's name to the supplied array `name` and returns `ch`.

When `ch` is a leading Unicode surrogate character and the optional `trail_ch` is a trailing surrogate character, the primitive combines them to form a Unicode character code, writes the name of the combined

character to *name*, and returns the combined character code. (The macros `IS_LEADING_SURROGATE()` and `IS_TRAILING_SURROGATE()` may be used to check if a character is a surrogate.

When *ch* is -1, the primitive looks up the Unicode character name given in *name* and returns its character number. If the name is unknown it returns -1 and puts an error message in *name*.

If none of these apply, it puts an error message in *name* and returns -1.

```
right_align_columns(char *pat)
```

The `right_align_columns()` subroutine locates all lines containing a match for the regular expression pattern *pat*. It notes the ending column of each match. (It assumes that *pat* occurs no more than one per line.)

Then, if some matches end at an earlier column than others, it adds indentation before each match as needed, so all matches will end at the same column.

```
columnize_buffer_text(int buf, int width, int margin)
```

The `columnize_buffer_text()` subroutine takes the lines in the buffer *buf* and reformats them into columns. It leaves a margin between columns of *margin* spaces, and chooses the number of columns so that the resulting buffer is at most *width* characters wide (unless an original line in the buffer is already wider than *width*).

```
do_buffer_to_hex(char *b, char transp[256], ?int flags)
```

The `do_buffer_to_hex()` primitive writes a hex view of the current buffer to the buffer *b*, creating or emptying it first. It ignores any narrowing in the original buffer. It uses the 256 character *transp* array to help construct the last column of the hex view; each character from the buffer will be replaced by the character at that offset in the *transp* array. If the buffer contains Unicode characters with codes higher than 255, they'll appear as-is.

If a buffer might contain Unicode characters, the primitive uses a display format that leaves room for 16 bits per character; otherwise it uses a format with room for 8 bits per character. The optional *flags* argument, if 1, forces 8 bits per character. If any character in the buffer doesn't fit in 8 bits, only its lower 8 bits will be shown in the hex listing.

10.1.12 Comparing

```
int compare_buffer_text(int buf1, int pos1,
                        int buf2, int pos2, int fold)
int buffers_identical(int a, int b)
```

The `compare_buffer_text()` primitive compares two buffers, specified by buffer numbers, starting at the given offsets within each. If *fold* is nonzero, Epsilon performs case-folding as in searching before comparing each character, using the case-folding rules of the current buffer. The primitive returns the number of characters that matched before the first mismatch.

The `buffers_identical()` subroutine checks to see if two buffers, specified by their buffer numbers, are identical. It returns nonzero if the buffers are identical, zero if they differ. If neither buffer exists, they're considered identical; if one exists, they're different.

```
do_uniq(int incl_uniq, int incl_dups, int talk)
buf_sort_and_uniq(int buf)
```

The `do_uniq()` subroutine defined in `uniq.e` goes through the current buffer comparing each line to the next, and deleting each line unless it meets certain conditions.

If `incl_uniq` is nonzero, lines that aren't immediately followed by an identical line will be preserved. If `incl_dups` is nonzero, the first copy of each line that is immediately followed by one or more identical lines will be preserved. (The duplicate lines that follow will always be deleted.)

If `talk` is nonzero, the subroutine will display status messages as it proceeds.

The `buf_sort_and_uniq()` subroutine sorts the specified buffer and discards duplicate lines in it, with no status messages.

```
do_compare_sorted(int b1, int b2, char *only1,
                  char *only2, char *both)
```

The `do_compare_sorted()` subroutine works like the `compare-sorted-windows` command, but lets you specify the two buffers to compare, and the names of the three result buffers. Any of the result buffer names may be `NULL`, and the subroutine won't generate data for that buffer.

```
int tokenize_lines(int buf1, int **lines1, int *len1,
                  int buf2, int **lines2, int *len2)
int lcs(int *lines1, int len1, int *lines2, int len2, char *outbuf)
```

These primitives help to compute a minimum set of differences between the lines of two buffers `buf1` and `buf2`. See the implementation of the `diff` command for an example of their use.

Call the `tokenize_lines()` primitive first. It begins by counting the lines in each buffer (placing the results in `len1` and `len2`). Then it uses the `realloc()` primitive to make room in the arrays passed by reference as `lines1` and `lines2`, which may be null at the start. Each array will have room for one token (unique integer) for each line of its buffer. (The arrays may be freed after calling `lcs()`, or reused in later calls.)

The `tokenize_lines()` primitive then fills in the arrays with unique tokens, chosen so that two lines will have the same token if and only if they're identical.

The `lcs()` primitive takes the resulting arrays and line counts, and writes a list of shared line ranges to the specified buffer, one per line, in ascending order. Each line range consists of a line number for the first buffer, a line number for the second (both 0-based) and a line count. For instance, a line "49 42 7" indicates that the seven lines starting at line 49 in the first buffer match the seven lines starting at line 42 in the second (counting lines from 0).

```
int lcs_char(int buf1, int from1, int to1,
             int buf2, int from2, int to2, char *outbuf)
```

The `lcs_char()` primitive is a character-oriented version of the `tokenize_lines()` and `lcs()` primitives described above. It compares ranges of characters in a pair of buffers.

It writes a list of shared character ranges to the specified buffer, one per line, in ascending order. Each character range consists of a character offset for the first buffer relative to `from1`, a character offset for the second buffer relative to `from2`, and a character count. For instance, a line "49 42 7" in the output buffer indicates that the seven characters in the range `from1 + 47` to `from1 + 47 + 7` in the first buffer match the seven characters in the range `from2 + 42` to `from2 + 42 + 7` in the second.

```
int phoneticize_lines(int dest, int len)
```

The `phoneticize_lines()` primitive quickly finds sound codes for a list of words. It goes through the current buffer line by line. Each line should contain a word; non-word characters will be ignored. For each line, it writes a corresponding line to the `dest` buffer with a phonetic code for that word, a string of letters designed so that two words with similar sounds will have the same phonetic code. (It currently uses the Metaphone algorithm for this purpose.) The `len` value indicates the maximum length of each phonetic code to be produced.

10.1.13 Managing Buffers

```
int create(char *buf)
char *bufnum_to_name(int bnum)
int name_to_bufnum(char *bname)
int zap(char *buf)
buf_zap(int bnum)
int change_buffer_name(char *newname)
```

The `create()` primitive makes a new buffer. It takes the name of the buffer to create. If the buffer already exists, nothing happens. In either case, it returns the buffer number of the buffer.

Some primitives let you specify a buffer by name; others let you specify a buffer by number. Epsilon tries never to reuse buffer numbers, so EEL functions can look a buffer up by its buffer number to see if a particular buffer still exists. Functions that accept a buffer number generally start with `buf_`.

Use the `bufnum_to_name()` primitive to convert from a buffer number to the buffer's name. If no such buffer exists, it returns a null pointer. The `name_to_bufnum()` primitive takes a buffer name, and gives you the corresponding buffer number. If no such buffer exists, it returns zero.

The `zap()` primitive creates a buffer if necessary, but empties it of all characters if the buffer already exists. So calling `zap()` always results in an empty buffer. The `zap()` primitive returns the buffer number of the buffer, whether or not it needed to create the buffer. The `buf_zap()` primitive works like `zap()`, except the former takes a buffer number instead of a buffer name, and signals an error if no buffer with that number exists. Unlike `zap()`, `buf_zap()` cannot create a buffer. Neither primitive switches to the emptied buffer.

The `change_buffer_name()` primitive renames the current buffer to the indicated name. If there is already a buffer with the new name, the primitive returns 0, otherwise the buffer is renamed and the primitive returns 1.

```
int exist(char *buf)
int buf_exist(int bnum)
delete_buffer(char *buf)
delete_user_buffer(char *buf)
buf_delete(int bnum)
drop_buffer(char *buf)      /* buffer.e */
char *temp_buf()           /* basic.e */
int tmp_buf()              /* basic.e */
```

The `exist()` primitive tells whether a buffer with a particular name exists. It returns 1 if the buffer exists, 0 if not. The `buf_exist()` primitive does the same thing, but takes a buffer number instead of a buffer name.

The `delete_buffer()` primitive removes a buffer with a given name. It also removes all windows associated with the buffer. The `buf_delete()` primitive does the same thing, but takes a buffer number. Epsilon signals an error if the buffer does not exist, if it contains a running process, or if one of the buffer's

windows could not be deleted. If the buffer might have syntax highlighting in it, use the `delete_user_buffer()` subroutine instead; it cleans up some data needed by syntax highlighting.

The `drop_buffer()` subroutine deletes the buffer, but queries the user first like the kill-buffer command if the buffer contains unsaved changes.

The EEL subroutine `temp_buf()`, defined in `basic.e`, uses the `exist()` primitive to create an unused name for a temporary buffer. It returns the name of the empty buffer it creates. The `tmp_buf()` subroutine creates a temporary buffer like `temp_buf()`, but returns its number instead of its name.

```
buffer char *bufname;
buffer int bufnum;
```

The `bufname` variable returns the name of the current buffer, and the `bufnum` variable gives its number. Setting either switches to a different buffer. If the indicated buffer does not exist, nothing happens. Use this method of switching buffers only to temporarily switch to a new buffer; use the `to_buffer()` or `to_buffer_num()` subroutines described on page 466 to change the buffer a window will display.

To set the `bufname` variable, use the syntax `bufname = new value;`. Don't use `strcpy()`, for example, to modify it.

```
int buffer_size(char *buf)
int buf_size(int bnum)
int get_buf_point(int buf)
set_buf_point(int buf, int pos)
```

The `buffer_size()` and `buf_size()` subroutines return the size in characters of the indicated buffer (specified by its name or number). The `get_buf_point()` subroutine returns the value of point in the indicated buffer. The `set_buf_point()` subroutine sets point in the specified buffer to the value `pos`. These are all defined in `buffer.e`.

10.1.14 Catching Buffer Changes

```
user buffer short call_on_modify;
on_modify() /* buffer.e */
zeroed buffer (*buffer_on_modify)();
buffer char _buf_readonly;
check_modify(int buf)
```

If the buffer-specific `call_on_modify` variable has a nonzero value in a particular buffer, whenever any primitive tries to modify that buffer, Epsilon calls the EEL subroutine `on_modify()` first. By default, that subroutine calls the `normal_on_modify()` subroutine, which aborts the modification if the buffer-specific variable `_buf_readonly` is nonzero, indicating a read-only buffer, and does various similar things.

But if the `buffer_on_modify` buffer-specific function pointer is nonzero for that buffer, `on_modify()` instead calls the subroutine it indicates. That subroutine may wish to call `normal_on_modify()` itself.

An `on_modify()` function can abort the modification or set variables. But if it plans to return, it must not create or delete buffers, or permanently switch buffers.

One of `normal_on_modify()`'s tasks is to handle read-only buffers. There are several types of these, distinguished by the value of the `_buf_readonly` variable, which if nonzero indicates the buffer is read-only. A value of 1 means the user explicitly set the buffer read-only. The value 2 means Epsilon automatically set the buffer read-only because its corresponding file was read-only.

A value of 3 indicates pager mode; this is just like a normal read-only buffer, but if the user action causing the attempt at buffer modification happens to be the result of the `<Space>` or `<Backspace>` keys, Epsilon cancels the modification and pages forward or backward, respectively. In other types of read-only buffers, this happens only if the `readonly-pages` variable permits it. The value 4 means the `all-read-only` variable caused that buffer to be set read-only.

The `check_modify()` primitive runs the `on_modify()` function on a specified buffer (if `call_on_modify` is nonzero in that buffer). You can use this if you plan to modify a buffer later but want any side effects to happen now. If the buffer is marked read-only, this function will abort with an error message. If the buffer is in virtual mode and its cursor is positioned in virtual space, Epsilon will insert whitespace characters to reach the virtual column. Because this can change the value of `point`, you should call `check_modify()` before passing the values of `spots` to any function.

For example, suppose you write a subroutine to replace the previous character with a '+', using a statement like `replace(point - 1, '+')`. Suppose `point` has the value 10, and appears at the end of a line containing 'abc' (in column 3). Using virtual mode, the user might have positioned the cursor to column 50, however. If you used the above statement, Epsilon would call `replace()` with the value 9. Before replacing, Epsilon would call `on_modify()`, which, in virtual mode, would insert tabs and spaces to reach column 50, and move `point` to the end of the inserted text. Then Epsilon would replace the character 'c' at buffer position 9 with '+'. If you call `check_modify(bufnum)`; first, however, Epsilon inserts its tabs and spaces to reach column 50, and `point - 1` correctly refers to the last space it inserted.

```
reset_modified_buffer_region(char *tag)
int modified_buffer_region(int *from, int *to, ?char *tag)
```

Sometimes an EEL function needs to know if a buffer has been modified since the last time it checked. Epsilon can maintain this information using tagged buffer modification regions.

An EEL function first tells Epsilon to begin collecting this information for the current buffer by calling the `reset_modified_buffer_region()` primitive and passing a unique tag name. (Epsilon's syntax highlighting uses a modified buffer region named `needs-color`, for instance.) Later it can call the `modified_buffer_region()` primitive, passing the same tag name. Epsilon will set its `from` and `to` parameters to indicate the range of the buffer that has been modified since the first call.

For example, say a buffer contains six characters `abcdef` when `reset_modified_buffer_region()` is called. Then the user inserts and deletes some characters resulting in `abxyf`. A `modified_buffer_region()` would now report that characters in the range 2 to 4 have been changed. If the buffer contains many disjoint changes, `from` will indicate the start of the first change, and `to` the end of the last.

The `modified_buffer_region()` primitive returns 0 if the buffer hasn't been modified since the last `reset_modified_buffer_region()` with that tag. In this case `from` and `to` will be equal. (They might also be equal if only deletion of text had occurred, but then the primitive wouldn't have returned 0.) It returns 1 if the buffer has been modified. If `reset_modified_buffer_region()` has never been used with the specified tag in the current buffer, it returns -1, and sets the `from` and `to` variables to indicate the whole buffer.

The tag may be omitted when calling `modified_buffer_region()`. In that case Epsilon uses an internal tag that's reset on each buffer display. So the primitive indicates which part of the current buffer has been modified since the last buffer display.

10.1.15 Listing Buffers

```
char *buffer_list(int start)
int buf_list(int offset, int mode)
```

The `buffer_list()` primitive gets the name of each buffer in turn. Each time you call this primitive, it returns the name of another buffer. It begins again when given a nonzero argument. When it has returned the names of all the buffers since the last call with a nonzero argument, it returns a null pointer.

The `buf_list()` primitive can return the number of each existing buffer, one at a time, like `buffer_list()`. The mode can be 0, 1, or 2, to position to the lowest-numbered buffer in the list, the last buffer returned by `buf_list()`, or the highest-numbered buffer, respectively. The `offset` lets you advance from these buffers to lower or higher-numbered buffers, by providing a negative or positive offset. Unlike `buffer_list()`, this primitive lets you back up or go through the list backwards.

For example, this code fragment displays the names of all buffers, one at a time, once forward and once backward:

```
s = buffer_list(1);
do {
    say("Forward %d: %s", name_to_bufnum(s), s);
} while (s = buffer_list(0));

i = buf_list(0, 2);
do {
    say("Back %d: %s", i, bufnum_to_name(i));
} while (i = buf_list(-1, 1));
say("Done.");
```

10.2 Display Primitives

10.2.1 Creating & Destroying Windows

```
window_kill()
window_one()
```

The `window_kill()` primitive removes the current window if possible, in the same way as the kill-window command does. The `window_one()` primitive eliminates all but the current window, as the command one-window does.

```
remove_window(int win)
```

The `remove_window()` primitive deletes a window by handle or number. If you delete a tiled window, Epsilon expands other windows as needed to fill its space. You cannot delete the last remaining tiled window.

```
int give_window_space(int dir)
#define BLEFT 0 /* direction codes */
#define BTOP 1
#define BRIGHT 2
#define BBOTTOM 3
```

The `give_window_space()` primitive deletes the current window. It expands adjacent windows in the specified direction into the newly available space, returning 0. If there are no windows in the specified direction, it does nothing and returns 1.

```

window_split(int orientation)
#define HORIZONTAL      (0)
#define VERTICAL        (1)

```

The `window_split()` primitive makes two windows from the current window, like the commands `split-window` and `split-window-vertically` do. The argument to `window_split()` tells whether to make the new windows appear one on top of the other (with argument `HORIZONTAL`) or side-by-side (with argument `VERTICAL`). The standard EEL header file, `eel.h`, defines the macros `HORIZONTAL` and `VERTICAL`. The primitive returns zero if it could not split the window, otherwise nonzero. When you split the window, Epsilon automatically remembers to call the `prepare_windows()` and `build_mode()` subroutines during the next redisplay.

```

user short window_handle;
user short window_number;
next_user_window(int dir)

```

You may refer to a window in two ways: by its *window handle* or by its *window number*.

Epsilon assigns a unique window handle to a window when it creates the window. This window handle stays with the window for the duration of that window's lifetime. To get the window handle of the current window, use the `window_handle` primitive.

The window number, on the other hand, denotes the window's current position in the window order. You can think of the window order as the position of a window in a list of windows. Initially the list has only one window. When you split a window, the two child windows replace it in the list. The top or left window comes before the bottom or right window. When you delete a window, that window leaves the list. The window in the upper left has window number 0. Pop-up windows always come after tiled windows in this order, with the most recently created (and therefore topmost) pop-up window last. The `window_number` primitive gives the window number of the current window.

Epsilon treats windows in a dialog much like pop-up windows, assigning each a window number and window handle. The stacking order of dialogs is independent of their window handles, however. Deleting all the windows on a dialog makes Epsilon remove the dialog. (Epsilon doesn't count windows with the `system_window` flag set when determining if you've deleted the last window.)

To change to a different window, you can set either the `window_handle` or `window_number` variables. Epsilon then makes the indicated window become the current window. Epsilon interprets `window_number` modulo the number of windows, so window number -1 refers to the last window.

Many primitives that require you to specify a window will accept either its handle or its number. Use `window_handle` to remember a particular window, since its number can change as you add or delete windows.

You can increment or decrement the `window_number` variable to cycle through the list of available windows. But it's usually better to use the `next_user_window()` subroutine, passing it 1 to go to the next window or -1 to go to the previous one. This will skip over system windows.

```

int number_of_windows()
int number_of_popups()
int number_of_user_windows()
int is_window(int win)
#define ISTILED          1
#define ISPOPUP          2

```


The `number_of_windows()` primitive returns the total number of windows, and the `number_of_popups()` primitive returns the number of pop-up windows. The `number_of_user_windows()` subroutine returns the total number of windows, excluding system windows.

The `is_window()` primitive accepts a window handle. It returns `ISTILED` if the value refers to a conventional tiled window, `ISPOPUP` if the value refers to a pop-up window or a window in a dialog, or 0 if the value does not refer to a window. Unlike most window functions, it accepts only a window handle, not a window number.

10.2.2 Window Resizing Primitives

```
user window short window_height;
user window short window_width;
int text_height()
int text_width()
int window_content_width()
```

The `window_height` variable contains the height of the current window in lines, including any mode line or borders. Setting it changes the size of the window. Each window must have at least one line of height. The `window_width` variable contains the width of the current window, counting any borders the window may have. If you set these variables to illegal values, Epsilon will adjust them to the closest legal values.

The `text_height()` and `text_width()` primitives, on the other hand, exclude borders and mode lines from their calculations, returning only the number of lines or columns of the window available for the display of text.

If the buffer has been set to display line numbers, `text_width()` doesn't count the columns used for them, but the similar `window_content_width()` primitive does. With line numbers off, the two return the same value.

```
int window_edge(int orien, int botright)
#define TOPLEFT      (0)
#define BOTTOMRIGHT  (1)
```

The `window_edge()` primitive tells you where on the screen the current window appears. For the first parameter, specify either `HORIZONTAL` or `VERTICAL`, to get the column or row, respectively. For the second parameter, provide either `TOPLEFT` or `BOTTOMRIGHT`, to specify the corner. Counting starts at the upper left corner of the screen, which has 0 for both coordinates.

10.2.3 Preserving Window Arrangements

```
struct window_info {
    short left, top, right, bottom;
    short textcolor, hbordcolor;
    short vbordcolor, titlecolor;
    short borders, other, bufnum;
    int point, dpoint;
    /* primitives fill in before this line */
    int dcolumn;
    short prevbuf;
};
```

```

get_window_info(int win, struct window_info *p)
low_window_info(int win, struct window_info *p)
window_create(int first, struct window_info *p)
low_window_create(int first, struct window_info *p)
select_low_window(int wnum, int top, int bot,
                  int lines, int cols)

```

Epsilon has several primitives that are useful for recording a particular window configuration and reconstructing it later.

The `get_window_info()` subroutine fills a structure with information on the specified window. The information includes the window's size and position, its selected colors, and so forth. It uses the `low_window_info()` primitive to collect some of the information, then fills in the rest itself by inspecting the window.

After calling `get_window_info()` on each tiled window (obtaining a series of structures, each holding information on one window), you can restore that window configuration using the `window_create()` subroutine. It takes a pointer to a structure that `get_window_info()` filled in, and a flag that must be nonzero if this is the first window in the new configuration. It uses the `low_window_create()` primitive to create the window. The `point` or `dpoint` members of the structure may be -1 when you call `window_create()` or `low_window_create()`, and Epsilon will provide default values for `point` and `window_start` in the new window, based on values stored with the buffer. The window-creating functions remain in the window they create, so you can modify its window-specific variables.

After a series of `window_create()`'s, you must use the `select_low_window()` primitive to switch to one of the created windows (specifying it by window number or handle, as usual).

Using `window_create()` directly modifies windows, and Epsilon doesn't check that the resulting window configuration is legal. For example, you can define a set of tiled windows that leave gaps on the screen, overlap, or extend past the screen borders. The result of creating an illegal window configuration is undefined.

The first time you call `window_create()`, pass it a nonzero flag, and Epsilon will (internally) delete all tiled windows, and create the first window. Then call `window_create()` again, as needed, to create the remaining windows (pass it a zero flag). Finally, you must call the `select_low_window()` primitive. Once you begin using `window_create()`, Epsilon will not be able to refresh the screen correctly until you call the `select_low_window()` primitive to exit window-creation. The `top` and `bot` parameters specify the new values of the `avoid-top-lines` and `avoid-bottom-lines` variables, and set the variables to the indicated values while finishing window creation. The `lines` and `cols` parameters specify the size of the screen that was used to construct the old window configuration. All windows defined using `low_window_create()` are based on that screen size. When you call `select_low_window()`, Epsilon resizes all the windows you've defined so that they fit the current screen size.

```

save_screen(struct screen_info *p)
restore_screen(struct screen_info *p)

```

The `save_screen()` subroutine saves Epsilon's window configuration in a `struct screen_info` structure. The first time you call this subroutine on an instance of the `screen_info` structure, make sure its `wins` member is zero. The `restore_screen()` subroutine restores Epsilon's window configuration from such a structure.

10.2.4 Pop-up Windows

```

int add_popup(column, row, width, height, border, bnum)
/* macros for defining a window's borders */

```

```

/* BORD(BTOP, BSINGLE) puts single line on top */
#define BLEFT  0
#define BTOP   1
#define BRIGHT 2
#define BBOTTOM 3
#define BNONE  0
#define BBLANK 1
#define BSINGLE 2
#define BDOUBLE 3
#define BORD(side, val)      (((val) & 3) << ((side) * 2))
#define GET_BORD(side, bord) ((bord >> (side * 2)) & 3)
#define LR_BORD(val)         (BORD(BLEFT, (val)) + BORD(BRIGHT, (val)))
#define TB_BORD(val)         (BORD(BTOP, (val)) + BORD(BBOTTOM, (val)))
#define ALL_BORD(val)        (LR_BORD(val) + TB_BORD(val))

```

The `add_popup()` primitive creates a new pop-up window. It accepts the column and row of the upper left corner of the new window, and the width and height of the window (including any borders). The border parameter contains a code saying what sort of borders the window should have, and the `bnum` parameter gives the buffer number of the buffer to display in the window. The primitive returns the handle of the new window, or `-1` if the specified buffer did not exist, so Epsilon couldn't create the window. If the pop-up window is to become part of a dialog (see page 576), its size, position and border will be determined by the dialog, not the values passed to `add_popup()`. The window created becomes the current window.

You can define the borders of a window using macros from `codes.h`. For each of the four sides, you can specify no border, a blank border, a border drawn with a single line, or a border drawn with a double line, using the codes `BNONE`, `BBLANK`, `BSINGLE`, or `BDOUBLE`, respectively. Specify the side to receive the border with the macros `BLEFT`, `BTOP`, `BRIGHT`, and `BBOTTOM`. You can make a specification for a given side using the `BORD()` macro, writing `BORD(BBOTTOM, BDOUBLE)` to put a double-line border at the bottom of the window. Add the specifications for each side to get the complete border code.

You can use other macros to simplify the border specification. Write `LR_BORD(BSINGLE) + TB_BORD(BDOUBLE)` to produce a window with single-line borders on the left and right, and double-line borders above and below. Write `ALL_BORD(BNONE)` for a window with no borders at all, and the most room for text.

You can use the `GET_BORD()` macro to extract (from a complete border code) the specification for one of its sides. For example, to find the border code for the left-side border of a window with a border value of `bval`, write `GET_BORD(BLEFT, bval)`. If the window has a double-line border on that side, the macro would yield `BDOUBLE`.

```
int window_at_coords(int row, int col, ?int screen)
```

The `window_at_coords()` primitive provides the handle of the topmost window at a given set of screen coordinates. The primitive returns `-1` if no window occupies that part of the screen. The screen number parameter can be zero or omitted to refer to the main screen, but it is usually a screen number from the `mouse_screen` primitive.

```
int window_to_screen(int win)
```

The `window_to_screen()` primitive takes a window handle and returns its screen number. Windows that are part of a dialog box have nonzero screen numbers; in this version other windows always have a screen number of zero.

```
int screen_to_window(int screen)
```

The `screen_to_window()` primitive takes a screen number, as returned in the variable `mouse_screen`, and returns the window handle associated with it. If the screen number is zero, there may be several windows associated with it; Epsilon will choose the first one. In this version of Epsilon, nonzero screen numbers uniquely specify a window. It returns -1 if no windows are associated with that screen number.

```
user window int window_left;
user window int window_top;
```

The `window_left` and `window_top` primitive variables provide screen coordinates for the current window. You can set the coordinates of a pop-up window to move the window around. Epsilon ignores attempts to set these variables in tiled windows.

10.2.5 Pop-up Window Subroutines

```
view_buffer(char *buf, int last) /* complete.e */
view_buf(int buf, int last) /* complete.e */
```

Several commands in Epsilon display information using the `view_buffer()` subroutine. It takes the name of a buffer and displays it page by page in a pop-up window. The `view_buf()` subroutine takes a buffer number and does the same. Both take a parameter `last` which says whether the command is displaying the buffer as its last action.

If `last` is nonzero, Epsilon will create the window and then return. Epsilon's main command loop will take care of displaying the pop-up window, scrolling through it, and removing it when the user's done examining it. If the user executes a command like `find-file` while the pop-up window is still on the screen, Epsilon will remove the pop-up and continue with the command.

If `last` is zero, the viewing subroutine will not return until the user has removed the pop-up window (by pressing `<Space>` or `Ctrl-g`, for example). The command can then continue with its processing. The user won't be able to execute a prompting command like `find-file` while the pop-up window is still on the screen.

```
view_linked_buf(int buf, int last, int (*linker)())
int linker(char *link) /* linker function prototype */
```

Epsilon uses a variation of `view_buf()` to display some online help. The variation adds support for simple hyperlinks. The user can select one of the links in a page of displayed text and follow it to go to another page, or potentially to perform any other action. The `view_linked_buf()` subroutine shows a buffer with links.

The links are delimited with a `Ctrl-a` character before and a `Ctrl-b` character after each link. Epsilon's non-Windows documentation file `edoc` is in this format. (See page 554.) The `view_linked_buf()` subroutine will modify the buffer it receives, removing and highlighting the links before displaying it.

When the user follows a link, Epsilon will call the function pointer `linker` passed as a parameter to `view_linked_buf()`. The `linker` function, which may have any name, will receive the link text as a parameter.

```

    /* space at sides of viewed popup */
    short _view_left = 2;
    short _view_top = 2;
    short _view_right = 2;
    short _view_bottom = 6;

    short _view_border = ALL_BORD(BSINGLE);
    char *_view_title;      /* title for viewed popup */
    int view_loop(int win)

```

By default, the above subroutines create a pop-up window with no title and a single-line border, almost filling the screen. The window begins two columns from the left border and stops two columns from the right, and extends two lines from the top of the screen to six lines from the bottom. You can alter any of these values by setting the variables `_view_title`, `_view_border`, `_view_left`, `_view_top`, `_view_right`, and `_view_bottom`. Preserve the original default value using the `save_var` keyword. For example, this code fragment shows a buffer in a narrow window near the right edge of the screen labeled “Results” (surrounding a title with spaces often makes it more attractive):

```

    save_var _view_left = 40;
    save_var _view_title = " Results ";
    save_var _view_border = ALL_BORD(BDOUBLE);
    view_buffer(buf, 1);

```

A command that displays a pop-up window may want more control over the creation and destruction of the pop-up window than `view_buf()` and similar subroutines provide. A command can instead create its pop-up window itself, and call `view_loop()` to handle user interaction. The `view_loop()` subroutine takes the handle of the pop-up window to work with. The pop-up window may be a part of a dialog. (See the `display_dialog_box()` primitive described on page 576.)

The `view_loop()` subroutine lets the user scroll around in the window and watches for an unrecognized key (an alphabetic key, for example) or a key that has a special meaning. It returns when the user presses one of these keys or when the user says to exit. By default, the user can scroll off either end of the buffer and this subroutine will return. Set the `paging-retains-view` variable nonzero to prevent this. The `view_loop()` subroutine returns an `INP_` code from `eel.h` to indicate which user action made it exit. See that file for more information. The function that called `view_loop()` may choose to call `view_loop()` again, or to destroy the pop-up window and continue.

```

    error_if_input(int abort) /* complete.e */
    remove_final_view() /* complete.e */

```

If the user is entering a response to some prompt and gives another command that also requires a response, Epsilon aborts the command to prevent confusion. Such commands should call `error_if_input()`, which will abort if necessary. The subroutine also removes a viewed buffer, as described above, by calling `remove_final_view()` if necessary. If its `abort` parameter is nonzero, it will attempt to abort the outer command as well, if aborting proves necessary.

10.2.6 Window Attributes

```

    int get_wattrib(int win, int code)
    set_wattrib(int win, int code, int val)
    /* use these codes with get_wattrib() & set_wattrib() */

```

```

#define BLEFT          0
#define BTOP           1
#define BRIGHT        2
#define BBOTTOM       3
#define PBORDERS       4
#define PHORIZBORDCOLOR 5
#define PVERTBORDCOLOR 6
#define PTEXTCOLOR     7
#define PTITLECOLOR    8

```

The `get_wattrib()` and `set_wattrib()` primitives let you examine and modify many of a window's attributes, such as its position, size, or color. The `win` parameter contains the handle or number of the window to modify, and the `code` parameter specifies a particular attribute.

For the `code` parameter, you can specify one of `BLEFT`, `BTOP`, `BRIGHT`, or `BBOTTOM`, to examine or change the window's size or position. They refer to the screen coordinate of the corresponding edge. You can use `PBORDERS` to specify a new border code (see the description of `add_popup()` above). Or you can set one of the window's colors: each window has a particular color class it uses for its normal text (outside of any highlighted regions), its horizontal borders, its vertical borders, and its title text. Use the macros `PTEXTCOLOR`, `PHORIZBORDCOLOR`, `PVERTBORDCOLOR`, and `PTITLECOLOR`, respectively, to refer to these. Set them using a color class expression. (See page 113.) For example, the statement

```
set_wattrib(win, PTEXTCOLOR, color_class viewed_text);
```

makes the text in window `win` appear in the color the user selected for `viewed_text`.

```

window char system_window;
window char invisible_window;

```

Setting the window-specific primitive variable `system_window` to a nonzero value designates the current window as a system window. The user commands that switch windows will skip over system windows. Setting the window-specific primitive variable `invisible_window` to a nonzero value makes a window whose text Epsilon won't display (although it will display the border, if the window has one). Epsilon won't modify the part of the screen that would ordinarily display the window's text.

10.2.7 Buffer Text in Windows

```

to_buffer(char *buf)          /* buffer.e */
to_buffer_num(int bnum)       /* buffer.e */
window short window_bufnum;
switch_to_buffer(int bnum)
int give_prev_buf()           /* buffer.e */
prev_forget_buf(int buf)      /* buffer.e */
to_another_buffer(char *buf)
tiled_only()                  /* window.e */
int in_bufed()                 /* bufed.e */
quit_bufed()                   /* bufed.e */

```

The `to_buffer()` subroutine defined in `buffer.e` connects the current window to the named buffer, while `to_buffer_num()` does the same, but takes a buffer number. Both work by setting the

`window_bufnum` variable, first remembering the previous buffer displayed in the window so the user can easily return to it. The `window_bufnum` variable stores the buffer number of the buffer displayed in the current window.

Both of these functions check the file date of the new buffer and warn the user if the buffer's file has been modified on disk. The `switch_to_buffer()` subroutine skips this checking.

The `give_prev_buf()` subroutine retrieves the saved buffer number of the previous buffer displayed in the current window. If the previous buffer has been deleted, or there is no previous buffer for this window, it returns the number of another recently-used buffer. If it can't find any suitable buffer, it returns 0. The `prev_forget_buf()` subroutine says that the indicated buffer should be removed from the list of previous buffers.

The `to_another_buffer()` subroutine makes sure that `buf` is not the current buffer. If it is, the subroutine switches the current window to a different buffer. This subroutine is useful when you're about to delete a buffer.

Sometimes the user may issue a command that switches buffers, while in a `bufed` pop-up window, or some other type of pop-up window. Issuing `to_buffer()` would switch the pop-up window to the new buffer, rather than the underlying window. Such commands should call the `tilde_only()` subroutine before switching buffers. This subroutine removes any `bufed` windows or other unwanted windows, and returns to the original tiled window. It calls the `quit_bufed()` subroutine to remove `bufed` windows. If it can't remove some pop-up windows, it tries to abort the command that created them. The `quit_bufed()` subroutine uses the `in_bufed()` subroutine to determine if the current window is a `bufed` window.

```
user window int window_start;
user window int window_end;
fix_window_start()      /* window.e */
```

The `window_start` variable provides the buffer position of the first character displayed in the current window. Epsilon's `redisplay` sets this variable, but you can also set it manually to change what part of the buffer appears in the window. When Epsilon updates the window after a command, it makes sure that point is still somewhere on the screen, using the new value for `window_start`. If not, it alters `window_start` so point is visible.

The `window_end` variable provides the buffer position of the last character displayed in the window. Epsilon's `redisplay` sets this variable. Setting it does nothing.

The `fix_window_start()` subroutine adjusts `window_start`, if necessary, so that it occurs at the beginning of a line.

```
int get_window_pos(int pos, int *row, int *col)
int window_line_to_position(int row)
```

The `get_window_pos()` function takes a buffer position and finds the window row and column that displays the character at that position. It puts the row and column in the locations that `row` and `col` point to. It returns 0 if it could find the position in the window, or a code saying why it could not.

A return value of 1 means that the position you gave doesn't appear in the window because it precedes the first position displayed in the window. If the given position doesn't appear in the window because it follows the last position displayed in the window, the function returns 2. A return value of 3 means that the position "appears" before the left edge of the screen (due to horizontal scrolling), and 4 means that the position "appears" too far to the right. It doesn't change the locations that `row` and `col` refer to when it returns 1 or 2.

The `window_line_to_position()` primitive takes the number of a row in the current window, and returns the buffer position of the first character displayed on that row. It returns -1 if the row number provided is negative or greater than the number of rows in the window.

```
user int line_in_window;
user int column_in_window;
```

The `line_in_window` and `column_in_window` primitives give you the position of point in the current window, as set by the last `refresh()`. Both variables start counting from 0. If you switch windows, Epsilon will not update these variables until the next `refresh()`.

```
int window_extra_lines()
build_window()
window_to_fit(int max)      /* window.e */
popup_near_window(int new, int old)
```

When buffer text doesn't reach to the bottom of a window, Epsilon blanks the rest of the window. The `window_extra_lines()` primitive gives the number of blank lines at the bottom of the window that don't correspond to any lines in the buffer.

Some of the functions that return information about the text displayed in a window only provide information as of the last redisplay. Due to buffer changes, their information may now be outdated. The `build_window()` primitive reconstructs the current window internally, updating Epsilon's idea of which lines of text go where in the window, how much will fit, and so forth. This primitive updates the value of `window_end`. It may also modify the `display_column` and `window_start` variables if displaying the window as they indicate doesn't get to point. The `build_window()` function also updates the values returned by the `window_line_to_position()`, `get_window_pos()`, and `window_extra_lines()` functions.

Use the `window_to_fit()` subroutine to ensure that a pop-up window is no taller than it needs to be. It sets the window's height so that it's just big enough to hold the buffer's text, but never more than `max` lines tall. The subroutine has no effect on windows that form part of a dialog.

The `popup_near_window()` subroutine tries to move a pop-up window on the screen so it's near another window. It also adjusts the height of the pop-up window based on its contents, by calling `window_to_fit()`. The `bufed` command uses this to position its pop-up buffer list near the tiled window from which you invoked it.

```
window_scroll(int lines)
```

The `window_scroll()` primitive scrolls the text of the current window up or down. It takes an argument saying how many lines up to scroll the current window. With a negative argument, this primitive scrolls the window down. (See page 472 for information on scrolling text left or right.)

10.2.8 Window Titles and Mode Lines

```
window_title(int win, int edge, int pos, char *title)
#define TITLECENTER      (0)
#define TITLELEFT(offset) (1 + (offset))
#define TITLERIGHT(offset) (-(1 + (offset)))
make_title(char *result, char *title, int room)
```


You can position a title on the top or bottom border of a window using the `window_title()` primitive. (Also see the `set_window_caption()` primitive described on page 578.) It takes the window number in `win` and the text to display in `title`. (It makes a copy of the text, so you don't need to make sure it stays around after your function returns.) The `edge` parameter must have the value of `BTOP` or `BBOTTOM`, depending on whether you want the title displayed on the top or bottom border of the window.

Construct the `pos` parameter using one of the macros `TITLELEFT()`, `TITLECENTER`, or `TITLERIGHT()`. The `TITLECENTER` macro centers the title in the window. The other two take a number which says how many characters away from the given border the title should appear. For example, `TITLERIGHT(3)` puts the title three characters away from the right-hand edge of the window.

Epsilon interprets the percent character “%” specially when it appears in the title of a window. Follow the percent character with a character from the following list, and Epsilon will substitute the indicated value for that sequence:

%c Epsilon substitutes the current column number, counting columns from 0.

%C Epsilon substitutes the current column number, counting columns from 1.

%d Epsilon substitutes the current display column, with a < before it, and a space after. However, if the display column has a value of 0 (meaning horizontal scrolling is enabled, but the window has not been scrolled), or -1 (meaning the window wraps long lines), Epsilon substitutes nothing.

%D Epsilon substitutes the current display column, but if the display column is -1, Epsilon substitutes nothing.

%l Epsilon substitutes the current line number.

%m Epsilon substitutes the text “ More ”, but only if characters exist past the end of the window. If the last character in the buffer appears in the window, Epsilon substitutes nothing.

%P Epsilon substitutes the percentage of point through the buffer, followed by a percent sign.

%p Epsilon substitutes the percentage of point through the buffer, followed by a percent sign. However, if the bottom of the buffer appears in the window, Epsilon displays Bot instead. Epsilon displays Top if the top of the buffer appears, and All if the entire buffer is visible.

%s Epsilon substitutes “* ” if the buffer's modified flag has a nonzero value, otherwise nothing.

%S Epsilon substitutes “*” if the buffer's modified flag has a nonzero value, otherwise nothing.

%h Epsilon substitutes the current hour in the range 1 to 12.

%H Epsilon substitutes the current hour in military time in the range 0 to 23.

%n Epsilon substitutes the current minute in the range 0 to 59.

%e Epsilon substitutes the current second in the range 0 to 59.

%a Epsilon substitutes “am” or “pm” as appropriate.

Note: For the current time, use a sequence like `%2h:%02n %a` for “3:45 pm” or `%02H:%02n:%02e` for “15:45:21”.

%% Epsilon substitutes a literal “%” character.

%< Indicates that redisplay may omit text to the left, if all of the information will not fit.

%> Puts any following text as far to the right as possible.

With any of the numeric sequences, you can include a printf-style field width specifier between the % and the letter. You can use the same kinds of field width specifiers as C's printf() function. In column 9, for example, the sequence %4c expands to " 9", %04c expands to "0009", and %-4c expands to "9 ".

You can expand title text in the same way as displaying it would, using the make_title() primitive. It takes the title to expand, a character array where it will put the resulting text, and a width in which the title must fit. It returns the actual length of the expanded text.

```
prepare_windows()      /* disp.e */
window char _window_flags;
#define FORCE_MODE_LINE 1
#define NO_MODE_LINE   2
#define WANT_MODE_LINE 4

build_mode()           /* disp.e */
assemble_mode_line(char *line) /* disp.e */
set_mode(char *mode)    /* disp.e */
buffer char *major_mode; /* EEL variable */
user char mode_format[60];
clean_mode(char *mode)
```

Whenever Epsilon thinks a window's mode line or title may be out of date, it arranges to call the prepare_windows() and build_mode() subroutines during the next redisplay. The prepare_windows() subroutine arranges for the correct sort of borders on each window. This sometimes depends on the presence of other windows. For example, tiled windows get a right-hand border only if there's another window to their right. This subroutine will be called before text is displayed.

By default, prepare_windows() puts a mode line on all tiled windows, but not on any pop-up windows. You can set flags in the window-specific _window_flags variable to change this. Set FORCE_MODE_LINE if you want to put a mode line on a pop-up window, or set NO_MODE_LINE to suppress a tiled window's mode line. The prepare_windows() subroutine interprets these flags, and alters the WANT_MODE_LINE flag to tell build_mode() whether or not to put a mode line on the window.

The build_mode() subroutine calls the assemble_mode_line() subroutine to construct a mode line, and then uses the window_title() primitive to install it.

The assemble_mode_line() subroutine calls the set_mode() subroutine to construct the part of the mode line between square brackets (the name of the current major mode and a list of minor modes).

While many changes to the mode line require a knowledge of EEL, you can do some simple customizations by setting the variable mode_format. Edit these variables with set-variable, using the percent character sequences listed above. For example, if you wanted each mode line to start with a line and column number, you could add the text " Line %l Col %c " to mode_format.

An EEL function can add text to the start of a particular buffer's mode line by setting the buffer-specific variable mode_extra. Call the set_mode_message() subroutine to do this. It takes a pointer to the new text, or NULL to remove the current buffer's extra text. Internet FTPs use this to display the percent of a file that's been received (and similar data).

The set_mode() subroutine gets the name of the major mode from the buffer-specific major_mode variable, and adds the names of minor modes after it.

You can add new minor modes by defining a function with a name that starts with show_minor_mode_. It must take one parameter, a character pointer. When called, it should copy the minor mode name to the character pointer if the mode is in effect.

Sometimes Epsilon constructs variable or function names that include the current mode's name, to permit a mode to define its own value for some function. For instance, saving a file looks for a variable named *modename-add-final-newline*, where *modename* is the current mode's name. Since the mode name may contain characters that aren't valid in a variable name, functions can call the `clean_mode()` subroutine to get a version of the major mode with any invalid characters removed. Only (lowercased) alphanumerics, `_` and `-` (converted to `_`) will be copied to mode.

```
display_more_msg(int win)
```

The `display_more_msg()` subroutine makes the bottom border of the window `win` display a "More" message when there are characters past the end of the window, by defining a window title that uses the `%m` sequence.

10.2.9 Normal Buffer Display

Epsilon provides many primitives for altering the screen contents. This section describes those relating to the automatic display of buffers that happens after each command, as described below.

```
refresh()
maybe_refresh()
```

The `refresh()` primitive does a standard screen refresh, showing the contents of all Epsilon windows. The `maybe_refresh()` primitive calls `refresh()` only if there is no type-ahead. This is usually preferred since it lets Epsilon catch up with the user's typing more quickly. Epsilon calls the latter primitive after each command executes.

```
user window char build_first;
user buffer char must_build_mode;
user char full_redraw;
user char all_must_build_mode;
```

Epsilon normally displays each window line by line, omitting lines that have not changed. When a command has moved point out of the window, Epsilon must reposition the display point (the buffer position at which to start displaying text) to return point to the window. However, Epsilon sometimes does not know that repositioning is required until it has displayed the entire window. When it discovers that point is not in the window, Epsilon moves the display point to a new position and immediately displays the window again. Certain commands which would often cause this annoying behavior set the `build_first` variable to prevent it.

When the `build_first` variable is set, the next redisplay constructs each window internally first, checks that point is in the window, and only then displays it. The variable is then set back to zero. A `build_first` redisplay is slower than a normal redisplay, but it never flashes an incorrect window.

Epsilon "precomputes" most of the text of each mode line, so it doesn't have to figure out what to write each time it updates the screen. Setting the `must_build_mode` variable to 1 warns Epsilon that any mode lines for the current buffer must be rebuilt. The `make_mode()` subroutine in `disp.e` sets this to 1, and Epsilon rebuilds the mode lines of all windows displaying this buffer.

Setting the `all_must_build_mode` variable to 1 is like setting `must_build_mode` to 1 for all buffers.

Setting the `full_redraw` variable rebuilds all mode lines, as well as any precomputed information Epsilon may have on window borders, screen colors, and so forth.

It is necessary to set `full_redraw` when two parameters affecting the display have been changed. Make the `full_redraw` variable nonzero if the size of the tab character has changed, or if the display class of any character has been changed via the `_display_class` array.

Each time the mode line changes, the `make_mode()` subroutine calls any function whose name starts with `do_when_make_mode_`. Such functions receive no parameters.

```
screen_messed()
```

The `screen_messed()` primitive causes the next `refresh()` to completely redraw the entire screen.

```
user window int display_column;
```

The window-specific variable `display_column` determines how Epsilon displays long lines. If negative, Epsilon displays buffer lines too big to fit on one screen line on multiple screen lines, with a `\` or graphic character (see the `_display_characters` variable described below) to indicate that the line has been wrapped. If `display_column` is 0 or positive, Epsilon only displays the part of a line that fits on the screen. Epsilon also skips over the initial `display-column` columns of each line when displayed. Horizontal scrolling works by adjusting the display column.

```
int next_screen_line(int n)
int prev_screen_line(int n)
```

The `next_screen_line()` primitive assumes point is at the beginning of a screen line, and finds the `n`th screen line following that one by counting columns. It returns the position of the start of that line.

The `prev_screen_line()` primitive is similar. It returns the start of the `n`th screen line before the one point would be on. It does not assume that point is at the start of a screen line.

If Epsilon is scrolling long lines of text rather than wrapping them (because `display_column` is greater than or equal to zero), these primitives go to the beginning of the appropriate line in the buffer, not the `display_column`'th column. In this mode, `next_screen_line(1)` is essentially the same as `nl_forward()`, and `prev_screen_line(0)` is like `to_begin_line()`.

Screen Dimensions

```
short screen_cols;
short screen_lines;
```

The `screen_cols` and `screen_lines` primitives contain the number of columns and lines on Epsilon's main display (the OS window's size, or the current terminal's size, as appropriate). They are set when Epsilon starts up, using values provided by the operating system (or, for the Windows version, by the registry). Don't set these variables directly. Use the `resize_screen()` primitive described below.

```
short want_cols;
short want_lines;
```

The `want_cols` and `want_lines` primitives contain the values the user specified through the `-vc` and `-vl` switches, respectively, described on page 17. If these variables are 0, it means the user did not explicitly specify the number of lines or columns to display.

```
term_init()           /* video.e */
term_cmd_line()       /* video.e */
term_mode(int active) /* video.e */
```

Epsilon's standard startup code calls the subroutine `term_init()` when you start Epsilon, and `term_cmd_line()` when it wants to change Epsilon's window size to match what the user specified on the command line. (It changes sizes based on the command line *after* it restores any saved session.) The `term_mode()` subroutine controls switching when you exit Epsilon or run a subprocess. Its argument is 1 when entering Epsilon again (when a `shell()` call returns, for example) and 0 when exiting.

```
resize_screen(int lines, int cols)
```

Use the `resize_screen()` primitive to tell Epsilon to display a different number of lines or columns. It scales all the windows to the new screen dimensions, and then sets the `screen_lines` and `screen_cols` variables to the new screen size.

Character Display

```
buffer char *_display_class;
user buffer short tab_size;
char *_echo_display_class;
```

Modifying the character array `_display_class` lets you alter the way Epsilon displays characters. There is one position in the array for each of the 65,536 possible characters in a buffer. The code at each position determines how Epsilon displays the character when it appears in a buffer. This code is a *display code*.

Epsilon lets each character occupy one or more screen positions. For example, the Control-A character is usually shown in two characters on the screen as “^A”. The number of columns the `<Tab>` character occupies depends on the column it starts in. Epsilon uses the display codes 0 through 6 to produce the various multi-character representations it is capable of, as described below.

Besides these multi-character display codes, Epsilon provides a way to have one character display as another. If the display code of a character is not one of the special display codes 0 through 7, Epsilon interprets the display code as a graphics character. This graphics character becomes the single-column representation.

For example, if the display code for 'A' is 'B' (that is, if the value of `_display_class['A']` is the character 'B'), wherever an 'A' appears in the buffer, a 'B' will appear on the screen when it is displayed. The character is still really an 'A', however: only searches for 'A' will find it, an 'A' will be written if you save the file, and so forth. This facility is especially useful for supporting national character sets.

If a display code is from 0 to 7, it has a special meaning. By default, all characters have such a display code. These numbers have been given names in the file `codes.h`, and we'll use the names in this discussion for clarity.

Epsilon displays a character with display code `BNORMAL` as the character itself. If character 65, the letter 'A', has display code `BNORMAL` it is the same as if it had display code 65.

Epsilon displays a character with display code `BTAB` as a tab. The character is displayed as the number of blanks necessary to reach the next tab stop. The buffer-specific primitive variable `tab-size` sets the number of columns from one tab stop to the next. By default its value is eight.

A character with display code `BNEWLINE` goes to the start of the next line when displayed, as newline does normally.

Epsilon displays a character with display code BC as a control character. It is displayed as the ^ character, followed by the original character exclusive-or'ed with 64, and with the high bit stripped. BM and BMC are similar, with the prefix being M- and M-^, respectively.

Epsilon displays a character with display code BHEX as a hexadecimal number in the form 'xB7'. Specifically, the representation has the letter 'x', then the two-character hexadecimal character code (or four digits for character codes over 255). You can change many of the characters Epsilon uses for its representations of newlines, tabs, hex characters, and so forth; see below.

Finally, Epsilon for Windows uses the display code BDOUBLEWIDTH for an experimental feature supporting those East Asian languages that have double-width characters in a fixed-width font. In this version the feature is normally disabled, and the display code is reserved for future use.

By default, the tab character has code BTAB, the newline character has code BNEWLINE, and the other control characters have code BC. Control characters with the eighth bit set have code BMC. All other characters have code BNORMAL.

The variable `_display_class` is actually a buffer-specific pointer to the array of display codes. Normally, all these pointers refer to the same array, contained in the variable `_std_disp_class` defined in `cmdline.e`. You can create other arrays if you wish to have different buffers display characters in different ways. Whenever you change the `_display_class` variable, `build_first` must be set to make the change take effect, as described above.

When displaying text in the echo area, Epsilon uses the display class array pointed to by the `_echo_display_class` variable. It can have the same values as `_display_class`.

```
char _display_characters[ ];
buffer char *buffer_display_characters;
```

It is possible to change the characters Epsilon uses to display certain parts of the screen such as the border between windows. Epsilon gets such characters from the `_display_characters` array. This array contains the line-drawing characters that form window borders, the characters Epsilon uses in some of the display modes set by `set-show-graphic`, and the characters it uses to construct the scroll bar. The `set-display-characters` command may be used to set these characters.

If the buffer-specific variable `buffer_display_characters` is non-null in a buffer, Epsilon uses it in place of the `_display_characters` variable whenever it displays that buffer. You can use this to provide a special window border, scroll bar, or similar for a particular buffer. Epsilon's `change-show-spaces` command uses this variable, too.

```
int expand_display(char *to, char *from)
```

The `expand_display()` primitive expands characters to the multicharacter representations they would have if displayed on the screen. It returns the length of the result.

Character Widths and Columns

```
int display_width(int ch, int col)
move_to_column(int col)
int column_to_pos(int col)
```

The number of characters that fit on each screen line depends on the display codes of the characters in the line. Epsilon moves characters with multi-character representations as a unit to the next screen line when they don't fit at the end of the previous one (except in horizontal scrolling mode). Tab characters also vary in

width depending upon the column they start in. There are several primitives that count screen columns using display class information.

The `display_width()` primitive is the simplest. It returns the width a character `ch` would have if it were at column `col`. The `move_to_column()` primitive moves to column `col` in the current line, or to the end of the line if it does not reach to column `col`. The `column_to_pos()` subroutine accepts a column number but doesn't move point; instead it returns the buffer position of that column.

```
int horizontal(int pos)
int current_column()
int get_column(int pos)      /* indent.e */
int get_indentation(int pos) /* indent.e */
int give_position_at_column(int p, int col)
to_column(int col)          /* indent.e */
indent_to_column(int col)   /* indent.e */
insert_to_column(int start, int end) /* indent.e */
indent_like_tab()          /* indent.e */
```

The `horizontal()` primitive returns the number of columns from point to position `pos`. Point doesn't change. It must be before `pos`. The primitive returns `-1` if there is a character of display code `BNEWLINE` between point and `pos`. This primitive assumes that point is in column 0.

The `current_column()` primitive uses the `horizontal()` primitive to return the number of the current column.

The `get_column()` subroutine returns the column number of a given buffer position. The `get_indentation()` subroutine returns the indentation of the line containing position `pos`.

The `give_position_at_column()` subroutine returns the buffer position of the character at column `col` on the line containing position `p`; the column may be `-1` to retrieve the position of the end of `p`'s line.

The `to_column()` subroutine indents so that the character immediately after point winds up in column `col`. It replaces any spaces and tabs before point with the new indentation. It doesn't modify any characters after point.

The `indent_to_column()` subroutine indents so that the next non-whitespace character on the line winds up in column `col`. It replaces any spaces and tabs before or after point. The older `insert_to_column()` function inserts spaces and tabs in a simpler way, and must be told what starting column to assume. In most cases `indent_to_column()` is a better choice.

The `indent_like_tab()` subroutine indents like inserting a `<Tab>` character at point would. However, it respects the `indent-with-tabs` variable and avoids using tabs when the variable is zero. It also converts spaces and tabs immediately before point so that they match `indent-with-tabs` and use the minimum number of characters.

```
force_to_column(int col) /* indent.e */
```

The `force_to_column()` subroutine tries to move to column `col`. If the line doesn't reach to that column, the function indents out to the column. If the column occurs inside a tab character, the function converts the tab to spaces.

```
user window short cursor_to_column;
to_virtual_column(int col) /* basic.e */
int virtual_column()      /* basic.e */
int virtual_mark_column() /* basic.e */
```

The window-specific `cursor_to_column` variable lets users position the cursor in a part of a window where there are no characters. It's normally `-1`, and the cursor stays on the character after point. If it's non-negative in the current window, Epsilon puts the cursor at the specified column in the window instead. Epsilon resets `cursor_to_column` to `-1` whenever the buffer changes, or point moves from where it was when you last set `cursor_to_column`. (Epsilon only checks these conditions when it redisplay the window, so you can safely move point temporarily.)

Similarly, the window-specific `mark_to_column` variable lets you position the mark in a part of a window where there are no characters. Epsilon uses this variable when it displays a region that runs to the mark's position, and swaps the variable with `cursor_to_column` when you exchange the point and mark. It's normally `-1`, so Epsilon highlights up to the actual buffer position of the mark. If it's non-negative in the current window, Epsilon highlights up to the specified column instead. Epsilon resets `mark_to_column` to `-1` just as described above for `cursor_to_column`.

The `to_virtual_column()` subroutine positions the cursor to column `col` on the current line. It tries to simply move to the correct position in the buffer, but if no buffer character begins at that column, it uses the `cursor_to_column` variable to get the cursor to the right place without immediately modifying the buffer. In that case, it also arranges for an `on_modify()` function to run just before any subsequent modification of that buffer. If the cursor is still in virtual space when the `on_modify()` function runs, that function adds spaces and tabs as needed to reach the given column, so any text inserted is positioned at the desired column. (EEL programmers usually use `move_to_column()` to position to a column instead of this function's more complicated semantics, and use `indent_to_column()` or `to_column()` to explicitly add whitespace characters when needed.)

The `virtual_column()` subroutine provides the column the cursor would appear in: either the value of the `cursor_to_column` variable, or (if it's negative) the current column. Similarly, the `virtual_mark_column()` subroutine provides the column for the mark, taking `mark_to_column` into account.

```
tab_convert(int from, int to, int totabs)
hack_tabs(int offset)
int maybe_indent_rigidly(int rev)
int standard_tab_cmd(int (*func)(), int indent, int flags)
```

The `tab_convert()` subroutine converts tabs to spaces in the specified region when its parameter `totabs` is zero. When `totabs` is nonzero, it converts spaces to tabs.

The `hack_tabs()` subroutine converts tabs to spaces in the `offset` columns following point. If `offset` is negative, the function converts tabs in the columns preceding point.

Commands bound to `<Tab>` often call the `maybe_indent_rigidly()` subroutine. If a region's been highlighted, this subroutine indents it using the `indent-rigidly` command and then returns nonzero. Otherwise, it returns zero. If its parameter `rev` is nonzero, the subroutine unindents; a command bound to `Shift-Tab` often provides a nonzero `rev`, but for commands on `<Tab>` this is typically zero.

The `standard_tab_cmd()` subroutine packages a lot of standard functionality often found on the `<Tab>` key for a mode. A mode's command for the `<Tab>` key can call it, passing it the mode's indenter function (see below).

If there's a highlighted region, it calls `maybe_indent_rigidly()` and returns 1. Otherwise, if the cursor is past the current line's indentation, it indents to the next level, where the `indent` parameter supplies the number of columns to indent for each level, returning 2. (Epsilon uses the `tab-size` if `indent` is zero or negative.)

Next, if the user's pressed `<Tab>` twice or more in a row, it indents one level more and returns 3; otherwise it calls the indenting function and returns 4. But if the indenting function supplies no indentation

(placing the line at the left margin), there was no indentation on the line before, and the 1 bit in the flags argument is provided, the function adds one level of indentation and returns 5.

```
buffer int (*indenter)(); /* EEL variable */
user buffer int auto_indent; /* EEL variable */
prev_indenter() /* indent.e */
zeroed int indenter_action; // Bits for why we're indenting.
#define IACT_AUTO_INDENT 1
#define IACT_AUTO_FILL 2
#define IACT_COMMENT 4
#define IACT_FILL_COMMENT 8
#define IACT_AUTO_FILL_COMMENT 16
#define IACT_REINDENT_PREV 32
```

The normal-character command provides a hook for automatic line indentation when it inserts the newline character. If the buffer-specific variable `auto-indent` is nonzero, the normal-character command will call the function pointed to by the variable `indenter`, a buffer-specific function pointer, after inserting a newline character. By default, it calls the `prev_indenter()` subroutine, which indents to the same indentation as the previous line.

Various other functions call a buffer's indenter function (if nonzero) at certain times.

Whenever Epsilon calls a buffer's indenter function, it sets the `indenter_action` variable to a value that indicates why it's indenting, so modes can vary indentation behavior based on context. The value `IACT_AUTO_INDENT` indicates the user pressed `(Enter)` in a buffer with auto-indenting enabled, as above.

The value `IACT_AUTO_FILL` indicates Epsilon broke a long line due to auto-filling, and it's now indenting the new line. The commenting commands `indent-for-comment` and `set-comment-column` use the value `IACT_COMMENT` when indenting comments.

The fill-comment command uses the value `IACT_FILL_COMMENT` when creating new lines for the comment. (In some modes it doesn't use the indenter when creating new lines.) Auto-filling of long comment lines uses the value `IACT_AUTO_FILL_COMMENT`.

The value `IACT_REINDENT_PREV` indicates the user pressed `(Enter)`, and a mode-specific setting caused Epsilon to reindent the existing line. See the `default-reindent-previous-line` variable.

The possible values of `indenter_action` are arranged as bits, so a function can use a bitmask to select particular sets of conditions.

10.2.10 Displaying Status Messages

```
int say(char *format, ...)
int sayput(char *format, ...)
```

The `say()` primitive displays text in the echo area. It takes a printf-style format string, and zero or more other parameters, as described on page 480. The `sayput()` primitive is similar, but it positions the cursor at the end of the string. Each returns the number of characters displayed.

```
int note(char *format, ...)
int noteput(char *format, ...)
int unseen_msgs()
int unseen_msgs_time()
wait_for_unseen_msgs()
drop_pending_says()
short expire_message;
```

When you use the `say()`, `sayput()`, or `error()` primitives (`error()`'s description appears on page 534) to display a message to the user, Epsilon ensures that it remains on the screen long enough for the user to see it (the `see-delay` variable controls just how long) by delaying future messages. Messages that must remain on the screen for a certain length of time are called *timed messages*.

The `note()` and `noteput()` primitives work like `say()` and `sayput()`, respectively, but their messages can be overwritten immediately. These untimed messages should be used for “status” messages that don’t need to last (“95% done”, for example).

Epsilon copies the text of each timed message to the `#messages#` buffer. It doesn’t copy untimed messages (but see the `show_text()` primitive below).

The `unseen_msgs()` primitive returns the number of unexpired timed messages. When the user presses a key, and there are unseen messages, Epsilon immediately displays the most recent message waiting to be displayed, and discards all pending timed messages.

The `unseen_msgs_time()` primitive returns the time remaining for the current timed message. It returns 0 if there are no timed messages, or -1 if the current timed message has an infinite delay (and thus won’t be replaced until the user presses a key).

The `wait_for_unseen_msgs()` subroutine delays until all timed messages have been displayed. Epsilon calls it before exiting. If one of the messages has an infinite delay set, it displays that message for several seconds and then returns regardless. Since only a key press can advance past a message with an infinite delay, it also reads and discards keys while displaying such messages.

The `drop_pending_says()` primitive makes Epsilon discard any timed messages that have not yet been displayed. It also makes the current message be untimed (as if it were generated by `note()`, not `say()`), so that the next `say()`, `note()`, or similar will appear immediately. It returns 0 if there were no timed messages, or 1 if there were (or the current message had not yet expired).

An EEL function sometimes needs to display some text in the echo area that is only valid until the user performs some action. For instance, a command that displays the number of characters in the buffer might wish to clear that count if the user inserts or deletes some characters. After displaying text with one of the primitives above, an EEL function may set the `expire_message` variable to 1 to tell Epsilon to clear that text on the next user key.

```
int aside(char *format, ...)
```

The `aside()` primitive copies text to the `#messages#` buffer like `say()` does, but without displaying anything on the screen.

```
int show_text(int column, int time, char *fmt, ...)
```

The `show_text()` primitive is the most general command for displaying text in the echo area. Like the other display primitives, it takes a printf-style format string, and returns the number of characters it displayed.

When Epsilon displays text in the echo area, you can tell it to begin at a particular column, and Epsilon will subdivide the echo area into two sections. You can then display different messages in each area independently of one another. When it’s necessary to display a very long message, Epsilon will combine the sections again and use the full display width. There are never more than two sections in the echo area.

In detail, the `show_text()` primitive tells Epsilon to begin displaying text in the echo area at the specified column, where the leftmost column is column 0. Epsilon then clears the rest of that echo area section, but doesn’t modify the other section.

Whenever you specify a column greater than zero in `show_text()`, Epsilon will subdivide the echo area at that column. It will clear any text to the right of the newly-displayed text, but not any text to its left.

Epsilon will recombine the sections of the echo area under two conditions: whenever you write text starting in column 0 that begins to overwrite the next section, and whenever you write the empty string "" at column 0. When Epsilon recombines sections, it erases the entire echo area before writing the new text.

Specifying a column of -1 acts just like specifying column 0, making Epsilon display the text at the left margin, but it also tells Epsilon to position the cursor right after the text.

The `time` says how long in hundredths of a second Epsilon must display the message before moving on and displaying the next message, if any. As with any timed message, when the user presses a key, Epsilon immediately displays the last message waiting, skipping through any pending messages. A value of 0 for `time` means the message doesn't have to remain for any fixed length of time. A value of -1 means that Epsilon may not go on to the next message until it receives a keystroke; such messages will never time out.

Most of the other echo area display primitives are equivalent to some form of `show_text()`, as shown in the following table:

<code>note("abc")</code>	<code>show_text(0, 0, "abc")</code>
<code>say("abc")</code>	<code>show_text(0, see_delay, "abc")</code>
<code>noteput("abc")</code>	<code>show_text(-1, 0, "abc")</code>
<code>sayput("abc")</code>	<code>show_text(-1, see_delay, "abc")</code>
<code>aside("abc")</code>	<code>show_text(-2, 1, "abc")</code>

Just as Epsilon copies timed messages created with `say()` or `sayput()` to the `#messages#` buffer, the text from a `show_text()` call will be copied if its `time` is nonzero. Epsilon treats a `time` of 1 (hundredth of a second) the same as zero (it's untimed), but still copies it to the `#messages#` buffer. A column of -2 has a special meaning: Epsilon copies the resulting text to the `#messages#` buffer if `time` is nonzero, but doesn't display it at all, just like the `aside()` primitive.

```
delayed_say(int flags, int before, int after, char *fmt, ...)
```

The `delayed_say()` primitive tells Epsilon to display some text later. It's intended for operations that may take some time. EEL code may display a message provisionally before starting some potentially lengthy operation, telling Epsilon to display it only after a certain amount of time has elapsed, and then cancel the pending message when it finishes. The message will only appear if the operation actually took a long time.

The `before` parameter says how long to wait, in hundredths of a second, before displaying the message, which is built using the `printf`-style format string `fmt` and any following parameters. If the resulting message is zero-length, it simply cancels any delayed say message that has not yet been displayed.

If the `before` time elapses and the message hasn't been canceled, Epsilon displays it, ensuring it isn't overwritten by a following message for at least `after` hundredths of a second. Epsilon saves the message in the `#messages#` buffer if `after` is nonzero. If `flags` is 1, Epsilon positions the cursor after the message; otherwise it doesn't.

```
int mention(char *format, ...)
user char mention_delay;
```

The `mention()` primitive acts like `sayput()`, but displays its string only after Epsilon has paused waiting for user input for `mention_delay` tenths of a second. It doesn't cause Epsilon to wait for input, it just arranges things so that if Epsilon does wait for input and the required delay elapses, the message is displayed and the wait continues. Writing to the echo area with `say()` or the like cancels any pending `mention()`. By default, `mention_delay` is 0.

```
int muldiv(int a, int b, int c)
```

The `muldiv()` primitive takes its arguments and returns the value $a * b / c$, performing this computation using 128-bit arithmetic. It's useful in such tasks as showing "percentage complete" while operating on a large buffer. Simply writing `point * 100 / size()` in EEL would use 64-bit arithmetic, as EEL always does, and on very large buffers the result could be wrong. (Previous versions of EEL used 32-bit arithmetic, where such overflows were more likely and this function was more useful.)

10.2.11 Printf-style Format Strings

Primitives like `say()` along with several others take a particular pattern of arguments. The first argument is required. It is a character pointer called the *format string*. The contents of the format string determine what other arguments are necessary.

Characters in the format string are copied to the echo area except where a percent character '%' appears. The percent begins a sequence which interpolates the value of an additional argument into the text that will appear in the echo area. The sequence has the following pattern, in which square brackets [] enclose optional items:

```
% [ ' ] [ - ] [ number ] [ . number ] character
```

In this pattern *number* may be either a string of digits or the character '*'. If the latter, the next argument provided to the primitive must be an int, and its value is used in place of the digits.

The meaning of the sequence depends on the final character:

- c** The next argument must be an int. (As explained previously, a character argument is changed to an int when a function is called, so it's fine here too.) The character with that ASCII code is inserted in the displayed text. For example, if the argument is 65 or 'A', the letter A appears, since the code for A is 65.
- d** The next argument must be an int. A sequence of characters for the decimal representation of that number is inserted in the displayed text. For example, if the argument is 65 the characters '6' and '5' are produced. With the ' modifier, values of more than four digits are shown with commas, as in 12,345,678.
- x** The next argument must be an int. A sequence of characters for the hexadecimal (base 16) representation of that number is inserted in the displayed text. For example, if the argument is 65 the characters '4' and '1' are produced (since the hexadecimal number 0x41 is equal to 65 in base 10). No minus sign appears with this representation.
- o** The next argument must be an int. A sequence of characters for the octal representation of that number is inserted in the displayed text. For example, if the argument is 65 the three characters "101" are produced (since the octal number 101 is equal to 65 in base 10). No minus sign appears with this representation.
- s** The next argument, which must be a string, is copied to the displayed text.
- q** The next argument, which must be a string, is copied to the displayed text, but quoted for inclusion in a regular expression. In other words, any characters from the original string that have a special meaning in regular expressions are copied with a percent character (%) before them. See page 66 for information on regular expressions.

- r** The next argument (which must be a string containing a file name in absolute form) is copied to the displayed text, after being converted to relative form. Epsilon calls the `relative()` primitive, described on page 510, to do this.
- f** The next argument must be a string, typically the name of a file. This sequence is just like `%s` except when used in a primitive that displays text in the echo area, such as `say()`, and the entire text to be displayed is too wide to fit in the available room. In that case, Epsilon calls the `abbreviate_file_name()` subroutine defined in `disp.e` to abbreviate the file name so the entire message fits in the available width. If the displayed message is also recorded in the `#messages#` buffer, where no width restriction applies, the unabbreviated form of the message will be used.
- k** The next argument must be a number. Epsilon interprets it as a key code or Unicode character code, and interpolates the name of that key or character. If a number starting with zero appears after the `%` character, Epsilon uses the short form of the key name, if any. See the `key_value()` primitive to convert in the opposite direction, converting text (in the short form only) back to a key's numeric code.
- p** The next argument must be a number. Epsilon interprets it as a color class, and any following text appears in that color. This only works on those primitives that insert text into a buffer; the numeric argument is ignored in a `sprintf()` or `say()` or similar.
- e** The next argument must be a number. Epsilon interprets it as an Epsilon error code, one that Epsilon might store in the `errno` variable or return directly from certain primitives, and interpolates text describing the error.

The first number, if present, is the width of the field the argument will be printed in. At least that many characters will be produced, and more if the argument will not fit in the given width. If no number is present, exactly as many characters as are required will be used.

The extra space will normally be put before the characters generated from the argument. If a minus sign is present before the first number, however, the space will be put at the end instead.

If the first number begins with the digit 0, the extra space will be filled with zeros instead of spaces. A minus sign before the first number is ignored in this case.

The second number, if present, is the maximum number of characters from the string that will be displayed. For example, each of these lines displays the text, "Just an example":

```
say("Just %.2s example", "another");

say("Just %.*s example", 7-5, "another");
```

It may be tempting to substitute a string variable for the first parameter of `say()`. For example, when writing a function that displays its argument `msg` and pauses, it may seem natural to write `say(msg);`. This will work fine unless `msg` contains a `'%'` character. In that case, you will probably get an error message. Use `say("%s", msg);` instead.

```
user char in_echo_area;
```

The `in_echo_area` variable controls whether the cursor is positioned at point in the buffer, or in the echo area at the bottom of the screen. The `sayput()` primitive sets this variable, `say()` resets it, and it is reset after each command.

10.2.12 Other Display Primitives

```

term_write(int col, int row, char *str, int count,
           int colorclass, int clear)
term_write_attr(int col, int row, int chartowrite,
                int attrtowrite)
term_clear()
term_position(int col, int row)

```

The following primitives provide low-level screen control. The `term_clear()` primitive clears the screen. The `term_position()` primitive positions the cursor to the indicated row and column. The `term_write()` primitive puts characters directly on the screen. It puts `count` characters from `str` on the screen at the row and column in the specified `colorclass`. If `clear` is nonzero, it clears the rest of the line. The `term_write_attr()` primitive writes a single character at the specified location on the screen. Unlike `term_write()`, which takes a color class, this primitive takes a raw foreground/background color attribute pair. This primitive does nothing in Epsilon for Windows or under the X11 windowing system. For all these primitives, row and col start at 0, and the coordinate 0,0 refers to the upper left corner of the screen. If a keyboard macro is running, the `term_` primitives are ignored.

```

fix_cursor()    /* EEL subr. */
user int normal_cursor;
user int overwrite_cursor;
user int virtual_insert_cursor;
user int virtual_overwrite_cursor;
#define CURSOR_SHAPE(top, bot)      ((top) * 1000 + (bot))
#define GUI_CURSOR_SHAPE(height, width, offset) \
    ((offset * 1000 + (height)) * 1000 + (width))
int cursor_shape;

```

During screen refresh, Epsilon calls the EEL subroutine `fix_cursor()` to set the shape of the cursor. The subroutine chooses one of four variables depending upon the current modes, and copies its value into the `cursor_shape` variable, which holds the current cursor shape code. The Windows and X11 versions set the `gui_cursor_shape` variable in a similar way, from a different set of four variables. All these variables use values constructed by the `GUI_CURSOR_SHAPE()` or `CURSOR_SHAPE()` macros. See page 111 for details on these variables.

```

windows_set_font(char *title, int fnt_code)

```

The `windows_set_font()` primitive displays a Windows font selection dialog, allowing the user to pick a different font. It takes two parameters. `Title` specifies the title of the dialog box to display. The `fnt_code` says whether to set Epsilon's main font (`FNT_SCREEN`), the font for printing (`FNT_PRINTER`), or the font for Epsilon's dialogs (`FNT_DIALOG`). It's only available in the Windows GUI version of Epsilon. See the `has_feature` variable.

```

int using_oem_font(int screen)
char using_new_font;

```

The `using_oem_font()` primitive returns a nonzero value if the specified screen's font uses the OEM character set, rather than the ANSI/Windows character set. It takes a screen number. This primitive always returns 0 under Unix. The primitive variable `using_new_font` will be nonzero whenever some screen's font has been changed since the end of the last screen refresh, or when a new screen has been created, for example by displaying a dialog.

10.2.13 Highlighted Regions

Epsilon can display portions of a buffer in a different color than the rest of the buffer. We call each such portion a region. The most familiar region is the one between point and mark. Epsilon defines this region automatically each time you create a new buffer. (Also see the description of character coloring on page 487.)

Epsilon can display a region in several ways. The most common method corresponds to the one you see when you set the mark (by typing Ctrl-@) and then move around: Epsilon highlights each of the characters between point and mark. If you use the mark-rectangle command on Ctrl-x # to define a rectangular region, the highlighting appears on all columns between point and mark, on all lines between point and mark. The pop-up windows of the completion facility illustrate a third type of highlighting, where complete lines appear highlighted. The header file `codes.h` defines these types of regions as (respectively) `REGNORM`, `REGRECT`, and `REGLINE`. Epsilon won't do any highlighting for a region that has type 0.

A fourth type of highlighting, `REGINCL`, is similar to `REGNORM`, but includes an additional character at the end of the region. If a `REGNORM` region runs between position 10 and position 20 in the buffer, Epsilon would highlight the 10 characters between the two positions. But if the region were a `REGINCL` region, it would include 11 characters: the characters at positions 10 and 20, and all the characters between.

```
int add_region(spot from, spot to, int color,
              int type, ?int handle)
remove_region(int handle)
int modify_region(int handle, int code, int val)
window char _highlight_control;
```

You can define new regions with `add_region()`. It takes a pair of spots, a color class expression such as `color_class highlight`, a region display type (as described above), and, optionally, a numeric “handle”. It returns a nonzero numeric handle which you can use to refer to the region later. You can provide the spots in either order, and you may give the same spot twice (for example, in conjunction with `REGLINE`, to always highlight a single line). See page 113 for basic information on color classes, and page 419 for details on the syntax of color class expressions).

When you omit the `handle` parameter to `add_region()` (or provide a `handle` of zero) `add_region()` assigns an unused handle to the new region. You can also provide the handle of an existing region, and `add_region()` will assign the same handle to the new region. Any changes you make to one region by using `modify_region()` will now apply to both, and a single `remove_region()` call will remove both. You can link any number of regions in the same buffer in this way. The special handle value 1 refers to the region between point and mark that Epsilon creates automatically.

The `remove_region()` primitive takes a region handle, and deletes all regions with that handle. The handle may belong to a region in another buffer. Epsilon signals an error if the handle doesn't refer to any region.

The `modify_region()` primitive retrieves or sets some of the attributes of one or more regions. It takes a region handle, a modify code (one of the MR... codes below), and a new value. If you provide a “new value” that's out of range (such as -2, out of range for all modify codes), Epsilon will not change that attribute of the region, but will simply return its value. If you provide a valid new value, Epsilon will set that attribute of the region, and will return its previous value.

When several regions share the same handle, it's possible they will have different attribute settings. In this case, which region's attribute Epsilon returns is undefined. If you specify a new value for an attribute, it will apply to all regions with that handle.

The modify code `MRCOLOR` may be used to get or change a region's color class. The modify code `MRTYPE` may be used to get or change a region's display type, such as `REGRECT`. The codes `MRSTART` and

MREND may be used to set the two spots of a region; however, Epsilon will not return the spot identifier for a region, but rather its current buffer position.

You can force a region's starting and ending positions to specific columns using the modify codes MRSTARTCOL and MRENDCOL. For example, if a region runs from point to mark, and you set its MRSTARTCOL to 3, the region will start at column 3 of whatever line point is on. A column setting of -1 here makes Epsilon use the actual column value of the spot; no column will be forced.

You can set up a region to be "controlled" by any numeric global variable. Epsilon will display the region only if the variable is nonzero. This is especially useful because the variable may be window-specific. Since regions are associated with buffers, this is needed so that a buffer displayed in two windows can have a region that appears in only one of them.

The standard region between point and mark is controlled by the window-specific character variable `_highlight_control`. By default, other regions are not controlled by any variable. The modify code MRCONTROL may be used with `modify_region()` to associate a controlling variable with a region. Provide the global variable's name table index (obtainable through `find_index()`) as the value to set.

A region may be given an auto-delete property using the MRAUTODEL macro. Pass a value of 1 to enable auto-deleting, 0 to disable it. When you delete an auto-deleting region, it automatically deletes the two spots assigned to it. By default, no region is auto-deleting. The spots used for an auto-deleting region should not be shared with other regions, or system spots like `point_spot` or `mark_spot`.

```
set_region_type()      /* disp.e */
int region_type()      /* disp.e */
highlight_on()         /* disp.e */
highlight_off()        /* disp.e */
int is_highlight_on()  /* disp.e */
```

Several subroutines let you conveniently control highlighting of the standard region between point and mark. To set the type of the region, call the subroutine `set_region_type()` with the region type code, one of REGNORM, REGRECT, REGLINE, or REGINCL. This doesn't automatically turn on highlighting. Call `highlight_on()` to turn on highlighting, or `highlight_off()` to turn it off.

The `region_type()` subroutine returns the type of the current region, whether or not it's currently highlighted. The `is_highlight_on()` subroutine returns the type of the current region, but only if it's highlighted. It returns 0 if highlighting is off.

There are several subroutines that help you write functions that work with different types of regions. If you've written a function that operates on the text of a normal Epsilon region, add the following lines at the beginning of your function to make it work with inclusive regions and line regions as well:

```
save_spot point, mark;
fix_region();
```

When the user has highlighted an inclusive or line region, the `fix_region()` subroutine will reposition point and mark to form a normal Epsilon region with the same characters. (For example, in the case of a line region, Epsilon moves point to the beginning of the line.) The function also swaps point and mark so that point comes first (or equals mark, if the region happens to be empty). This is often convenient.

This procedure assumes your function doesn't plan to modify point or mark, just the characters between them, and it makes sure that point and mark remain in the same place. If your function needs to reposition the point or mark, try omitting the `save_spot` line. Your function will be responsible for determining where the point and mark wind up.

A function needs to do more work to operate on rectangular regions. If it's built to operate on all the characters in a region, without regard to rectangles or columns, the simplest approach may be to extract the rectangle into a temporary buffer, modify it there, and then replace the rectangle in the original buffer. Several Epsilon subroutines help you do this. For a concrete example, let's look at the function `fill_rectangle()`, defined in `format.e`. The fill-region command calls this function when the current region is rectangular.

```
// Fill paragraphs in rectangle between point and mark
// to marg columns (relative to rectangle's width if <=0).
fill_rectangle(marg)
{
    int width, orig = bufnum, b = tmp_buf();

    width = extract_rectangle(b, 0);
    save_var bufnum = b;
    mark = 0;
    margin_right = marg + (marg <= 0 ? width : 0);
    do_fill_region();
    xfer_rectangle(orig, width, 1);
    buf_delete(b);
}
```

The function begins by allocating a temporary buffer using `tmp_buf()`. Then it calls the `extract_rectangle()` subroutine to copy the rectangle into the temporary buffer. This function returns the width of the rectangle it copied. The call from `fill_rectangle()` passes the destination buffer number as the first parameter. Then `fill_rectangle()` switches to the temporary buffer and reformats the text. Finally, the subroutine copies the text back into its rectangle by calling `xfer_rectangle()` and deletes the temporary buffer. If the operation you want to perform on the text in the rectangle depends on any buffer-specific variables, be sure to copy them to the temporary buffer.

Now let's look at the two rectangle-manipulating subroutines `fill_rectangle()` calls in more detail.

```
extract_rectangle(int copybuf, int remove)
```

The `extract_rectangle()` subroutine operates on the region between point and mark in the current buffer. It treats the region as a rectangle, whether or not `region_type()` returns `REGRECT`. It can perform several different actions, depending upon its parameters. If `copybuf` is nonzero, the subroutine inserts a copy of the rectangle into the buffer with that buffer number. The buffer must already exist.

If `remove` is 1, the subroutine deletes the characters inside the rectangle. If `remove` is 2, the subroutine replaces the characters with spaces. If `remove` is 0, the subroutine doesn't change the original rectangle.

The subroutine always leaves point at the upper left corner of the rectangle and mark at the lower right. It return the width of the rectangle.

```
xfer_rectangle(int dest, int width, int overwrite)
```

The `xfer_rectangle()` subroutine inserts the current buffer as a rectangle of the given width into buffer number `dest`, starting at `dest`'s current point. If `overwrite` is nonzero, the subroutine copies on top of any existing columns. Otherwise it inserts new columns. In the destination buffer, it leaves point at the top left corner of the new rectangle, and mark at the bottom right. The point remains at the same position in the original buffer.

```
rectangle_standardize()
```

Functions that manipulate rectangles can sometimes use the `rectangle_standardize()` subroutine to simplify their logic. In a rectangular region, point may be at any one of the four corners of the rectangle. This subroutine moves point and mark so they indicate the same region, but with point at the lower right and mark at the upper left. It's like the rectangular region equivalent of the `fix_region()` subroutine.

```
do_shift_selects()
```

Commands bound to cursor keys typically select text when you hold down the shift key. They do this by calling `do_shift_selects()` as they start. This routine looks at the current state of the shift key and whether or not highlighting is already on, and turns highlighting on or off as needed, possibly setting point.

```
make_line_highlight()    /* complete.e */
remove_line_highlight()  /* complete.e */
```

The `make_line_highlight()` subroutine uses the `add_region()` primitive to create a region that highlights the current line of the current buffer. When Epsilon puts up a menu of options, it uses this function to keep the current line highlighted. The `remove_line_highlight()` subroutine gets rid of such highlighting.

10.2.14 Character Coloring

You can set the color of individual characters using the `set_character_color()` primitive. At first glance, this feature may seem similar to Epsilon's mechanism for defining highlighted regions. Both let you specify a range of characters and a color to display them with. But each has its own advantages.

Region highlighting can highlight the text in different ways: as a rectangle, expanded to entire lines, and so forth, while character coloring has no similar options. You can define a highlighted region that moves around with the point, the mark, or any other spot. Character coloring always remains with the characters.

But when there are many colored regions, using character coloring is much faster than creating a corresponding set of highlighted regions. If you define more than a few dozen highlighted regions, Epsilon's screen refreshes will begin to slow down. Character coloring, on the other hand, is designed to be very fast, even when there are thousands of colored areas. Character coloring is also easier to use for many tasks, since it doesn't require the programmer to allocate spots to delimit the ends of the colored region, or delete them when the region is no longer needed.

One more difference is the way you remove the coloring. For highlighted regions, you can turn off the coloring temporarily by calling `modify_region()`, or eliminate the region entirely by calling `remove_region()`. To do either of these, you must supply the region's handle, a value returned when the region was first created. On the other hand, to remove character coloring, you can simply set the desired range of characters to the special color -1. A program using character coloring doesn't need to store a series of handles to remove or modify the coloring.

Epsilon's code coloring functions are built on top of the character coloring primitives described in this section. See the next section for information on the higher-level functions that make code coloring work.

```
set_character_color(int pos1, int pos2, int color)
```

The `set_character_color()` primitive makes Epsilon display characters between `pos1` and `pos2` using the specified color class. Epsilon discards any previous color settings of characters in that range.

A color class of `-1` means the text will be “uncolored”. To display uncolored text, Epsilon uses the standard color class `text`. When a buffer is first created, every character is uncolored.

When you insert text in a buffer, it takes on the color of the character immediately after it, or in the case of the last character in the buffer, the character immediately before it. Characters inserted in an empty buffer are initially uncolored. Copying text from one buffer to another does not automatically transfer the color; Epsilon treats the new characters the same as any other inserted text. You can use the `buf_xfer_colors()` subroutine to copy text from one buffer to another and retain its coloring. See page 441.

Epsilon maintains the character colors set by this primitive independently of the highlighted regions created by `add_region()`. The `modify_region()` primitive will never change what `get_character_color()` returns, and similarly the `set_character_color()` primitive never changes the attributes of a region you create with `add_region()`. When Epsilon displays text, it combines information from both sources to determine the final color of each character.

When displaying a buffer, Epsilon uses the following procedure when determining which color class to use for a character:

- Make a list of all old-style highlighted regions that contain the character, and the color classes used for each.
- Add the character’s color as set by `set_character_color()` to this list.
- Remove color classes of `-1` from the list.

Next, Epsilon chooses a color class from the list:

- If the list of color classes is empty, use the `text` color class.
- Otherwise, if the list contains the `highlight` color class, use that.
- Otherwise, use the color class from the old-style highlighted region with the highest region number. If there are no old-style highlighted regions in the list, the list must contain only one color class, so use that.
- Finally, if we wound up selecting the `text` color class, and the `text_color` variable isn’t equal to `color_class text`, use the color class in the `text_color` variable instead of the `color_class text`.

Notice that when a region using the `highlight` color class overlaps another region, the `highlight` color class takes precedence.

```
buf_set_character_color(int buf, int from, int to, int color)
```

The `buf_set_character_color()` subroutine is a convenience function. It simply runs `set_character_color()` in the specified buffer `buf`, passing it the remaining parameters.

```
short get_character_color(int pos, ?int *startp, ?int *endp)
```

The `get_character_color()` primitive returns the color class for the character at the specified buffer position, as set by `set_character_color()`, or -1 if the character is uncolored, and will be displayed using the window's default color class.

You can also use the primitive to determine the extent of a range of characters all in the same color. If the optional pointer parameters `startp` and `endp` are non-null, Epsilon fills in the locations they point to with buffer positions. These specify the largest region of the buffer containing characters the same color as the one at `pos`, and including `pos`. For example, if the buffer contains a five-character word that has been colored blue, the buffer is otherwise uncolored, and `pos` refers to the second character in the word, then Epsilon will set `*startp` to `pos - 1` and `*endp` to `pos + 4`.

```
set_tagged_region(char *tag, int from, int to, short val)
short get_tagged_region(char *tag, int pos, ?int *from, int *to)
```

The character coloring primitives above are actually built from a more general facility that allows you to associate a set of attributes with a buffer range.

Each set of attributes consists of a tag (a unique string like "my-tag") and, for each character in the buffer, a number that represents the attribute. Each buffer has its own set of tags, and each tag has its own list of attributes, one for each character. (Epsilon stores the numbers in a way that's efficient when many adjacent characters have the same number, but nothing prevents each character from having a different attribute.)

The `set_tagged_region()` primitive sets the attribute of the characters in the range `from` to `to`, for the specified tag.

The `get_tagged_region()` primitive gets the attribute of the character at position `pos` in the buffer. If you provide pointers `from` and `to`, Epsilon will fill these in to indicate the largest range of characters adjacent to `pos` that have the same attribute as `pos`. Characters whose attributes have never been set for a given tag will have the attribute -1.

Epsilon's character color primitives `set_character_color()` and `get_character_color()` use a built-in tagged region with a tag name of "colors".

10.2.15 Code Coloring Internals

Epsilon's code coloring routines use the character coloring primitives above to do code coloring for various languages like C, TeX, and HTML. There are some general purpose code coloring functions that manage code coloring and decide what sections of a buffer need to be colored. Then, for each language, there are functions that know how to color text in that language.

The general purpose section maintains information on what parts of each buffer have already been colored. It divides each buffer into sections that are already correctly colored, and sections that may not be correctly colored. When the buffer changes, it moves its divisions so that the modified text is no longer marked "correctly colored". Whenever Epsilon displays part of a buffer, this part of code coloring recolors sections of the buffer as needed, and marks them so they won't be colored again unless the buffer changes. Epsilon only displays the buffer after the appropriate section has been correctly colored. This part also arranges to color additional sections of the buffer whenever Epsilon is idle, until the buffer has been completely colored.

The other part of code coloring does the actual coloring of C, TeX, and HTML buffers. You can write new EEL functions to tell Epsilon how to color other languages, and use the code coloring package's mechanisms for remembering which parts of the buffer have already been colored, and which need to be recolored. This section describes how to do this. (Also see page 585.)

```

buffer int (*recolor_range)();
    // how to color part of this buffer
buffer int (*recolor_from_here)();
    // how to find a good starting pos
int color_c_range(int from, int to)
    // how to color part of C buffer
int color_c_from_here(int safe)
    // how to find starting pos in C buffer
buffer char coloring_flags;
#define COLOR_DO_COLORING          1
#define COLOR_IN_PROGRESS          2
#define COLOR_MINIMAL              4
#define COLOR_INVALIDATE_FORWARD  8
#define COLOR_INVALIDATE_BACKWARD 16
#define COLOR_INVALIDATE_RESETS   32
#define COLOR_RETAIN_NARROWING    64
#define COLOR_IGNORE_INDENT       128
int must_color_through;

```

You must first write two functions and make the buffer-specific function pointers refer to them, in each buffer you want to color. For C/C++/EEL buffers, the c-mode command takes care of setting the function pointers. It also contains the lines

```

if (want_code_coloring)
    when_setting_want_code_coloring();

```

to actually turn on code coloring for the buffer if necessary.

The first function, which must be stored in the buffer-specific `recolor_range` variable, does the actual coloring of a part of the buffer. It takes two parameters `from` and `to` specifying the range of the buffer that needs coloring. It colors at least the specified range, but it may go past `to` and color more of the buffer. It returns the buffer position it reached, indicating that all characters between `from` and its return value are now correctly colored. In C buffers, the `recolor_range` function is named `color_c_range()`.

The `recolor_range` function may decide to mark some characters in the range “uncolored”, by calling `set_character_color()` with a color class of `-1`. Or it may assign particular color classes to all parts of the range to be colored. But either way, it should make sure all characters in the given range are correctly colored. Typically, a function begins by setting all characters between `from` and `to` to a default color class, then searching for elements which should be colored differently. Be sure that if you extend the range past `to`, you color all the characters between `to` and your new stopping point.

Epsilon remembers which parts of the buffer require coloring by using a tagged region (see page 488) named “needs-color”. A coloring routine may decide, while parsing a buffer, that some later or earlier section of the buffer requires coloring; if so, it can set the `needs-color` attribute of that section to `-1` to indicate this, and Epsilon will recolor that section of the buffer the next time it’s needed. Or it can declare that some other section of the buffer is already properly colored by setting that section’s attribute to `0`. It may also decide to examine the `must_color_through` variable, a buffer position marking the end of the region that really requires coloring right now. (Ordinarily, Epsilon expands this region to include entire color blocks.)

When the buffer’s modified, some of its coloring becomes invalid, and must be recomputed the next time it’s needed. Normally Epsilon invalidates a few lines surrounding the changed section. Some language modes tell Epsilon to automatically invalidate more of the buffer by setting flags in the buffer-specific

coloring_flags variable. (Other flags in this variable aren't normally set by language modes; code coloring uses them for bookkeeping purposes.)

COLOR_INVALIDATE_FORWARD indicates that after the user modifies a buffer, any syntax highlighting information after the modified region should be invalidated. COLOR_INVALIDATE_BACKWARD indicates that syntax highlighting information before the modified region should be invalidated.

COLOR_INVALIDATE_RESETS tells Epsilon that whenever it invalidates syntax highlighting in a region, it should also set the color of all text in that region to the default of -1. COLOR_RETAIN_NARROWING indicates that coloring should respect any narrowing in effect (instead of looking outside the narrowed area to parse the buffer in its entirety). COLOR_IGNORE_INDENT says that a simple change of indentation shouldn't cause any recoloring. Languages with no column-related highlighting rules may set this for better performance.

For many languages, starting to color at an arbitrary place in the buffer requires a lot of unnecessary work. For example, the C language has comments that can span many lines. A coloring function must know whether it's inside a comment before it can begin coloring. Similarly, a coloring function that began looking from the third character in the C identifier `id37` might decide that it had seen a numeric constant, and incorrectly color the buffer.

To simplify this problem, the coloring routines ensure that coloring begins at a safe place. We call a buffer position *safe* if the code coloring function can color the buffer beginning at that point, without looking at any earlier characters in the buffer.

When Epsilon calls the function in `recolor_range`, the value of `from` is always safe. Epsilon expects the function's return value to be safe as well; it must be OK to continue coloring from that point. For C, this means the returned value must not lie inside a comment, a keyword, or any other lexical unit. Moreover, inside the colored region, any boundary between characters set to different color classes must be safe. If the colored region contains a keyword, for example, Epsilon assumes it can begin recoloring from the start of that keyword. (If this isn't true for a particular language, its coloring function can examine the buffer itself to determine where to begin coloring.)

When Epsilon needs to color more of the buffer, it generally starts from a known safe place: either a value returned by the buffer's `recolor_range` function, or a boundary between characters of different colors. But when Epsilon first begins working on a part of the buffer that hasn't been colored before, it must determine a safe starting point. The second function you must provide, stored in the `recolor_from_here` buffer-specific function pointer, picks a new starting point. In C buffers, the `recolor_from_here` function is named `color_c_from_here()`.

The buffer's `recolor_from_here` function looks backward from point for a safe position and returns it. This may involve a search back to the start of the buffer. If Epsilon knows of a safe position before point in the buffer, it passes this as the parameter `safe`. (If not, Epsilon passes 0, which is always safe.) The function should respect the value of the `color-look-back` variable to limit searching on slow machines.

Epsilon provides two standard `recolor_from_here` functions that coloring extensions can use. The `recolor_by_lines()` subroutine is good for buffers where coloring is line-based, such as dired buffers. In such buffers the coloring needed for a line doesn't depend at all on the contents of previous lines. The `recolor_from_top()` subroutine has just the opposite effect; it forces Epsilon to start from the beginning of the buffer (or an already-colored place). This may be all that's needed if a mode's coloring function is very simple and quick.

Epsilon runs the code coloring functions while it's refreshing the screen, so running the EEL debugger on code coloring functions is difficult, since the debugger itself needs to refresh the screen. The best way to debug such functions is to test them out by calling them explicitly, using test-bed functions like these:

```
command debug_color_region()
{
```

```

        fix_region();
        set_character_color(point, mark, color_class default);
        point = color_algol_range(point, mark);
    }

    command debug_from_here()
    {
        point = color_algol_from_here(point);
    }

```

The first command above tries to recolor the current region, and moves past the region it actually colored. It begins by marking the region with a distinctive color (using the default color class), to help catch missing coloring. The second command helps you test your `from_here` function. It moves point backwards to the nearest safe position. Once you're satisfied that your new code-coloring functions work correctly, you can then set the `recolor_range` and `recolor_from_here` variables to refer to them.

To help find bugs in a code-coloring function, you can have Epsilon check whether it has colored every character in the range. This requires two steps. First, define `DEBUG_COLORING` in your EEL file before it includes `eel.h`. (This makes it intercept all your `set_character_color()` calls and send them to a debugging version.) Second, set the `want-debug-coloring` variable nonzero.

Bits in the `want-debug-coloring` variable enable features. The 1 bit makes Epsilon simply check if a coloring function colored every character. That is, if it was given a range `start...end` and returned the value `ret` (which must be `>= end`), it checks whether `set_character_color()` was called to assign a color on every character in the range `start...ret`. If it skips some, Epsilon will display a message and add a report to a debug buffer whose name is the original buffer's name with "debug-color-" prefixed. The 2 bit makes color debugging apply a distinctive color class to all the missed characters.

The 4 bit makes the debug buffer's report list not only every region that was missed, but also every region that had a color applied to it, and the 8 bit generates such a report every time your coloring function is invoked. Finally, the 16 bit has the debug buffer list every single call to `set_character_color()` by your coloring function. (These bits can be helpful in simply seeing what regions Epsilon is asking for, and what your function does in response.)

Detecting uncolored-region bugs only works on those coloring subroutines that are designed to color every character in their given range (plus any additional range they report they've colored, when they return a value greater than the range `start...end` they were given). But some modes' coloring functions (for example, HTML mode) keep track of which buffer regions require coloring in a different way, and other modes intentionally only apply colors to certain parts of the buffer. (For example, process mode uses this method only for user input, and colors process output as it arrives, based on any ANSI escape sequences it contains.) Color debugging isn't useful for spotting bugs in such modes, but its logging ability can be useful.

```

buffer int (*when_displaying)();
recolor_partial_code(int from, int to)
char first_window_refresh;
add_buffer_when_displaying(int buf, int (*func)())
delete_buffer_when_displaying(int buf, int (*func)())
default_when_displaying(int from, int to)
drop_all_colored_regions()
drop_coloring(int buf)

```

Epsilon calls the EEL subroutine pointed to by the buffer-specific function pointer `when_displaying` as it displays a window on the screen. It calls this subroutine once for each window, after determining which part of the buffer will be displayed, but before putting text for that window on the screen.

Epsilon sets the `first_window_refresh` variable prior to calling the `when_displaying` subroutine to indicate whether or not this is the first time a particular buffer has been displayed during a particular screen refresh. When a buffer appears in more than one window, Epsilon sets this variable to 1 before calling the `when_displaying` subroutine during the display of the first window, and sets it to zero before calling that subroutine during the display of the remaining windows. Epsilon sets the variable to 1 if the buffer only appears in one window. The value is valid only during a call to the buffer's `when_displaying` subroutine.

In a buffer with code coloring turned on, the `when_displaying` variable points to a subroutine named `recolor_partial_code()`. Epsilon passes two values to the subroutine that specify the range of the buffer that was modified since the last time the buffer was displayed. The standard `recolor_partial_code()` subroutine provided with Epsilon uses this information to discard any saved coloring data for the modified region of the buffer in the data structures it maintains. It then calls the two language-specific subroutines described at the beginning of this section as needed to color parts of the buffer.

You can tell Epsilon to run a function at display time by calling the `add_buffer_when_displaying()` subroutine. It arranges for the specified function to be called after code coloring has been done when displaying any window showing the specified buffer. The function will be called with no parameters. The `delete_buffer_when_displaying()` removes the specified function from that buffer's list of functions to be called at display time.

The `recolor_partial_code()` subroutine calls the `default_when_displaying()` function, which calls each such function set by `add_buffer_when_displaying()`. In most buffers without code coloring turned on, the `when_displaying` variable points to the `default_when_displaying()` function directly. Other functions assigned to `when_displaying` should call `default_when_displaying()` too.

The `drop_all_colored_regions()` subroutine discards coloring information collected for the current buffer. The next time Epsilon needs to display the buffer, it will begin coloring the buffer again. The `drop_coloring()` subroutine is similar, but lets you specify the buffer number. It also discards some data structures, so it's more suitable when the buffer is about to be deleted.

10.2.16 Colors

```
user int selected_color_scheme;
short _our_mono_scheme;
short _our_color_scheme;
short _our_gui_scheme;
short _our_unixconsole_scheme;
short *get_color_scheme_variable()
window short window_color_scheme;
buffer short buffer_color_scheme;
```

Epsilon stores color choices in *color scheme* variables. A color scheme specifies the color combination to use for each defined color class.

Epsilon's standard color schemes are defined in the file `stdcolor.e`. See page 425 for the syntax of color definitions. You can also create additional color schemes without loading an EEL file by using the `new_variable()` primitive, providing `NT_COLSCHEME` as the second parameter. Epsilon stores color schemes in its name table, just like variables and commands, so a color scheme may not have the same name as a variable or other name table entry. (Color classes, on the other hand, have their own unique "name space".)

The `selected_color_scheme` primitive variable contains the name table index of the color scheme to use. Setting it changes the current color scheme. Each time Epsilon starts up, it sets this variable from one of four other variables: `_our_gui_scheme` under Epsilon for Windows or in Epsilon for Unix under X11,

`_our_unixconsole_scheme` if Epsilon for Unix is running in console mode and the `USE_DEFAULT_COLORS` environment variable is set (so that Epsilon inherits the colors of its xterm), `_our_mono_scheme` if Epsilon is running on a monochrome display, or `_our_color_scheme` otherwise. When you use `set-color` to select a different color scheme, Epsilon sets one of these variables, as well as `selected_color_scheme`. The `get_color_scheme_variable()` subroutine returns a pointer to one of these variables, the one containing a color scheme index that's appropriate for the current environment. By default, these four variables refer to the color schemes `standard-gui`, `xterm-color`, `standard-mono` and `standard-color`, respectively.

If the window-specific variable `window_color_scheme` is nonzero in a window, Epsilon uses its value in place of the `selected_color_scheme` variable when displaying that window. Epsilon uses this when displaying borderless windows, so that each window has an entirely different set of color class settings. Also see the variable `text_color`.

Similarly, if the buffer-specific variable `buffer_color_scheme` is nonzero in a buffer, Epsilon uses its value instead of either `window_color_scheme` or `selected_color_scheme`.

```
user char monochrome;
```

The `monochrome` variable is nonzero if Epsilon believes it is running on a monochrome display. Epsilon tries to determine this automatically, but the `-vmono` and `-vcolor` flags override this. See page 17.

```
set_color_pair(int colorclass, int fg, int bg, ?int scheme)
int get_foreground_color(int colorclass, ?int raw)
int get_background_color(int colorclass, ?int raw)
```

The `set_color_pair()` primitive lets you set the colors to use for a particular color class within the current color scheme (or, if the optional `scheme` argument is nonzero, in that scheme). The first parameter is a `color_class` expression (see page 419); the next two parameters are 32-bit numbers that specify the precise color to use. Use the `MAKE_RGB()` macro to construct suitable numbers. See page 425.

The `get_foreground_color()` and `get_background_color()` primitives let you retrieve the colors specified for a given color class. Normally they return a specific foreground or background color, after Epsilon has applied its rules for defaulting color specifications. (See page 425.) Specify a nonzero `raw` parameter, and Epsilon will return the color class's actual setting. It may include one of the bits `ETRANSSPARENT`, `ECOLOR_COPY`, or `ECOLOR_UNKNOWN`.

The `ETRANSSPARENT` macro is a special code that may be used in place of a background color. It tells Epsilon to substitute the background color of the "text" color class in the current color scheme. You can also use it for a foreground color, and Epsilon will substitute the foreground color of the "text" color class.

The `ECOLOR_UNKNOWN` macro in a foreground color indicates there's no color information in the current scheme for the specified color class.

The `ECOLOR_COPY` macro in a foreground color tells Epsilon that one color class is to borrow the settings of another. The index of other color class replaces the color in the lower bits of the value; use the `COLOR_STRIP_ATTR()` macro to extract it.

Regardless of whether the `raw` parameter is used, a retrieved foreground color may include any of the font style bits `EFONT_BOLD`, `EFONT_UNDERLINED`, or `EFONT_ITALIC`.

When Epsilon looks up the foreground and background settings of a color class, it uses this algorithm.

First it checks if the foreground color contains the `ECOLOR_UNKNOWN` code. If so, it tries to retrieve first a class-specific default, and then a scheme-specific default. First it looks for that color class in the "color-defaults" color scheme. This scheme is where Epsilon records all color class specifications that

are declared outside any particular color scheme. If a particular color pair is specified as a default for that class, Epsilon uses that. If the color class has no default, Epsilon switches to the color class named "default" in the original color scheme and repeats the process.

Either the default setting for the color class or the original setting for the color class may use the `ECOLOR_COPY` macro. If so, then Epsilon switches to the indicated color class and repeats the above process. In the event that it detects a loop of color class cross-references or otherwise can't resolve the colors, it picks default colors.

Finally, if the resulting foreground or background colors use the `ETRANSSPARENT` bit, Epsilon substitutes the foreground or background color from the "text" color class.

```
int define_color_class(char *name, b32 fg, b32 bg)
```

The `define_color_class()` primitive creates a new color class. If `name` is the name of an existing color class, it simply sets its colors to the specified foreground and background pair, like `set_color_pair()`. Otherwise, it creates a new color class by that name, initialized to the specified colors. If `name` is `NULL` or "", it creates an anonymous color class. The resulting color class always uses the specified foreground and background colors, and the `set-color` command can't be used to customize it.

```
int alter_color(int colorclass, int color)
int rgb_to_attr(int rgb)
int attr_to_rgb(int attr)
```

The `alter_color()` primitive is an older way to set colors. When the argument `color` is `-1`, Epsilon simply returns the color value for the specified color class. Any other value makes the color class use that color. Epsilon then returns the previous color for that color class. (In Epsilon for Windows or under the X11 windowing system, this function will return color codes, but ignores attempts to set colors. Use `set_color_pair()` to do this.)

The colors themselves (the second parameter to `alter_color()`) are specified numerically. Each number contains a foreground color, a background color, and an indication of whether blinking or extra-bright characters are desired.

The `alter_color()` function uses 4-bit color attributes to represent colors. The foreground color is stored in the low-order 4 bits of the 8-bit color attribute, and the background color is in the high-order 4 bits. Epsilon uses a pair of 32-bit numbers to represent colors internally, so `alter_color()` converts between the two representations as needed.

The functions `rgb_to_attr()` and `attr_to_rgb()` can be used to perform the same conversion. The `rgb_to_attr()` function takes a 32-bit RGB value and finds the nearest 4-bit attribute, using Epsilon's simple internal rules, while `attr_to_rgb()` converts in the other direction.

```
int orig_screen_color()
```

In some environments, Epsilon records the original color attribute of the screen before writing text to it. The `orig_screen_color()` primitive returns this color code. If the `restore-color-on-exit` variable is nonzero, Epsilon sets the color class it uses after you exit (`color_class_after_exiting`) to this color. See page 113.

```
int number_of_color_classes()
char *name_color_class(int colclass)
```

The `number_of_color_classes()` primitive returns the number of defined color classes. The `name_color_class()` primitive takes the numeric code of a color class (numbered from 0 to `number_of_color_classes() - 1`) and gives the name. For example, if the expression `color_class mode_line` has the value 3, then the expression `name_color_class(3)` gives the string "mode-line". It returns NULL for anonymous color classes created by `define_color_class()`.

Each window on the screen can use different color classes for its text, its borders, and its titles (if any). When a normal, tiled window is created, Epsilon sets its color selections from the color classes named `text`, `horiz_border`, `vert_border`, and `mode_line`. When Epsilon creates a pop-up window, it sets the window's color selections from the color classes `text`, `popup_border`, and `popup_title`. See page 113 for a description of the other predefined color classes.

```
user window int text_color;
```

The `text_color` primitive contains the color class of normal text in the current window. You can get and set the other color classes for a window using the functions `get_wattrib()` and `set_wattrib()`.

10.3 File Primitives

10.3.1 Reading Files

```
int file_read(char *file, int transl)
```

The `file_read()` primitive reads the named file into the current buffer, replacing the text that was there. It returns an error code if an error occurred, or 0 if the read was successful. The `transl` parameter specifies the line translation to be done on the file. The buffer's `translation-type` variable will be set to its value. If `transl` is `FILETYPE_AUTO`, Epsilon will examine the file as it's read and set `translation-type` to an appropriate translation type.

```
int new_file_read(char *name, int transl,
                  struct file_info *f_info,
                  int start, int max)
```

The `new_file_read()` primitive reads a file like `file_read()` but provides more options. The `f_info` parameter is a pointer to a structure, which Epsilon fills in with information on the file's write date, file type, and so forth. The structure has the same format as the `check_file()` primitive uses (see page 505). If the `f_info` parameter is null, Epsilon doesn't get such information.

When Epsilon reads the file, it starts at offset `start` and reads at most `max` characters. You can use this to read only part of a big file. If `start` or `max` are negative, they are (individually) ignored: Epsilon starts at the beginning, or reads the whole file, respectively. The `start` parameter refers to the file before Epsilon strips `(Return)`'s, while `max` counts the characters after stripping.

```
int do_file_read(char *s, int transl) /* files.e */
buffer char _read_aborted;
int read_file(char *file, int transl) /* files.e */
int find_remote_file(char *file, int transl)
file_convert_read(int flags)
do_readonly_warning()
update_readonly_warning(struct file_info *p)
```

Instead of calling the above primitives directly, extensions typically call one of several subroutines, all defined in `files.e`, that do things beyond simply reading in the file. Each takes the same two parameters as `file_read()`, and returns either 0 or an error code.

The `do_file_read()` subroutine records the file's date and time, so Epsilon can later warn the user that a file's been modified on disk, if necessary. If the user aborted reading the file, `do_file_read()` sets the `_read_aborted` variable to 1; it uses the value 2 if an error occurred reading the file. Epsilon then warns the user if he tries to save the partial file. This subroutine also handles reading URLs by calling the `find_remote_file()` subroutine, and character set translations such as OEM translations (see page 500) by calling `file_convert_read()`.

The `read_file()` subroutine calls `do_file_read()`, then displays either an error message, if a read error occurred, or the message "New file." It also handles calling `do_readonly_warning()` when it detects a read-only file, or `update_readonly_warning()` otherwise. (The latter can turn off a buffer's read-only attribute, if the file is no longer read-only.)

```
int find_in_other_buf(char *file, int transl) /* files.e */
call_mode(char *file)                      /* files.e */
```

The `find_in_other_buf()` subroutine makes up a unique buffer name for the file, based on its name, and then calls `read_file()`. It then goes into the appropriate mode for the file, based on the file's extension, by calling the `call_mode()` subroutine. (See page 84.)

```
int find_it(char *fname, int transl)        /* files.e */
int std_find_it(char *fname)                /* files.e */
int ask_find_it(char *fname)                /* files.e */
int get_default_translation_type(char *fname) /* files.e */
int look_file(char *fname)                  /* buffer.e */
```

The `find_it()` subroutine first looks in all existing buffers for the named file, just as the `find-file` command would. If it finds the file, it simply switches to that buffer. (It also checks the copy of the file on disk, and warns the user if it's been modified.) If the file isn't already in a buffer, it calls `find_in_other_buf()`, and returns 0 or its error code. The `find_it()` subroutine uses the `look_file()` subroutine to search through existing buffers for the file.

While `find_it()` requires you to pass the appropriate translation type, the `std_find_it()` and `ask_find_it()` subroutines supply this themselves. The `std_find_it()` subroutine always uses the default translation type for a file with the given name. The `ask_find_it()` subroutine usually does the same, but if the calling command was invoked with a numeric prefix argument, it prompts the user for the translation rules. The `get_default_translation_type()` subroutine returns the default translation type for a given file name.

The `look_file()` subroutine, defined in `buffer.e`, returns 0 if no buffer has the file. Otherwise, it returns 1 and switches to the buffer by setting `bufnum`.

```
int do_find(char *file, int transl)          /* files.e */
```

Finally, the `do_find()` subroutine is at the top of this tree of file-reading functions. It checks to see if its "file name" parameter is a directory. If it is (or if it's a file pattern with wildcard characters), it calls `dired_one()` to run `dired` on the pattern. If it's a normal file, `do_find()` calls `find_it()`.

```
int err_file_read(char *file, int transl)    /* files.e */
```

Use the `err_file_read()` subroutine when you want to read a file that must exist, but you don't want all the extras that higher-level functions provide: checking file dates, choosing a buffer, setting up for read-only files, and so forth. It calls `file_read()` to read the file into the current buffer, and displays an error message if the file couldn't be read for any reason. It returns the error code, or 0 if there were no errors.

```
short abort_file_io;
```

By default, primitives that read and write files respond to the user pressing the abort key by asking whether they should abort the input/output operation. An EEL program can select a different behavior by using `save_var` to set the primitive variable `abort_file_io`. The default setting, `ABORT_ASK`, asks the user whether to abort the operation. If he says no, the operation continues. If he says yes, the primitive returns an error code, `EREADABORT` for reading primitives or `EWRITEABORT` for writing primitives. The setting `ABORT_ERROR` omits asking the user; it immediately returns an error code if the user aborts. The setting `ABORT_IGNORE` tells Epsilon to ignore the abort key and continue. The setting `ABORT_JUMP` makes pressing the abort key abort the current function by calling the `check_abort()` primitive, again without prompting first. (See page 534.)

10.3.2 Writing Files

```
int file_write(char *file, int transl)
```

The `file_write()` primitive attempts to write the current buffer to the named file. It returns 0 if the write was successful, or an error code if an error occurred. The `transl` parameter specifies the line translation to be done while writing the file. See the description of translation-type below.

```
int new_file_write(char *name, int transl,
                  struct file_info *f_info,
                  int start, int max)
```

```
#define FILE_IO_ATSTART      -1
#define FILE_IO_NEWFILE     -2
#define FILE_IO_TEMPFILE    -3
char *get_tempfile_name()
```

The `new_file_write()` primitive writes a file, like `file_write()`, but provides more options. The `f_info` parameter is a pointer to a structure, which Epsilon fills in with information on the file's write date, file type, and so forth, just after it finishes writing the file. The structure has the same format as the `check_file()` primitive uses (see page 505). If the `f_info` parameter is null, Epsilon doesn't gather such information.

Different values for the `start` parameter change Epsilon's behavior. In the usual case, pass the value `FILE_IO_ATSTART`. That makes Epsilon open or create the file normally and start writing at its beginning, replacing its existing contents.

A start value of `FILE_IO_NEWFILE` makes the `new_file_write()` call fail if the file already exists. You can set the `file_write_newfile` variable nonzero when calling the higher-level writing functions described below to ensure that `new_file_write()` receives this value.

A start value of `FILE_IO_TEMPFILE` makes Epsilon ignore the specified file name and pick an unused file name for a temporary file, in a directory designated for temporary files. The `get_tempfile_name()` primitive returns a pointer to the most recent temporary file created in this way (in a static buffer that will be reused for the next temporary file name).

With any of the above `start` codes, whatever Epsilon writes replaces the previous contents of the file. If `start` is greater than or equal to zero, though, Epsilon only rewrites a section of the existing file, starting at the offset specified by `start`, and the rest of the file's data will not change.

If the `max` parameter is non-negative, Epsilon writes only the specified number of characters. (Epsilon counts the characters before adding any `<Return>` characters or performing any encoding; the count is of characters in the current buffer.) If `max` is negative, Epsilon writes the entire buffer to the file.

The file-writing primitives use the `abort-file-io` variable to decide what to do if the user presses the abort key; see page 497.

```
int do_save_file(int backup, int checkdate,
                int getdate) /* files.e */
```

The `do_save_file()` subroutine saves the current buffer like the `save-file` command, but lets you skip some of the things `save-file` does. Set the `backup` parameter to 0 if you don't want a backup file created, even if `want-backups` is nonzero. Set `checkdate` to 0 if you don't want Epsilon to check that the file on disk is unchanged since it was read. Set `getdate` to 0 if you don't want Epsilon to update its notion of the file's date, after the file has been written.

The function returns 0 if the write was successful, 1 if an error occurred, or 2 if the function asked the user to confirm a questionable write, and the user decided not to write the file after all.

```
int ask_save_buffer()
int warn_existing_file(char *s)
```

A command can call the `ask_save_buffer()` subroutine before deleting a buffer with unsaved changes. It asks the user if the buffer should be saved before it's deleted, and returns non-zero if the user asked that the buffer be saved. The caller is responsible for actually saving the file.

Before writing to a user-specified file, a command may call the `warn_existing_file()` subroutine. This will check if the file already exists and warn the user that it will be overwritten. The subroutine returns zero if the file didn't exist, or if the user said to go ahead and overwrite it, or nonzero if the user said not to overwrite it.

10.3.3 Line Translation

Epsilon normally deals with files with lines separated by the `<Newline>` character. Windows, DOS and OS/2, however, generally separate one line from the next with a `<Return>` character followed by a `<Newline>` character. For this reason, Epsilon normally removes all `<Return>` characters from a file when it's read from disk, and places a `<Return>` character before each `<Newline>` character when a buffer is written to disk, in these environments. But Epsilon has several other line translation methods:

The `FILETYPE_BINARY` translation type tells Epsilon not to modify the file at all when reading or writing.

The `FILETYPE_MSDOS` translation type tells Epsilon to remove `<Return>` characters when reading a file, and insert a `<Return>` character before each `<Newline>` when writing a file.

The `FILETYPE_UNIX` translation type tells Epsilon not to modify the file at all when reading or writing. It's similar to `FILETYPE_BINARY` (but Epsilon copies buffer text to the system clipboard in a different way).

The `FILETYPE_MAC` translation type tells Epsilon to convert `<Return>` characters to `<Newline>` characters when reading a file, and to convert `<Newline>` characters to `<Return>` characters when writing a file.

The `FILETYPE_AUTO` translation type tells Epsilon to examine the contents of a file as it's read, and determine the proper translation type using a heuristic. Epsilon then reads the file using that translation type, and sets `translation-type` to the new value. Normally this value is only used when reading a file, not when writing one. If you try to write a file and specify a translation type of `FILETYPE_AUTO`, it will behave the same as `FILETYPE_MSDOS` (except in Epsilon for Unix, where it's the same as `FILETYPE_UNIX`).

```

user buffer int translation_type; /* EEL variable */
#define FORCED_TRANS              (8)
#define MAKE_TRANSLATE(e, t, f)  (((e) << 4) | ((t) & 7) \
                                  | ((f) ? FORCED_TRANS : 0))

#define GET_ENCODING(t)          ((t) >> 4)
#define GET_LINE_TRANSLATE(t)    ((t) & 7)
#define SET_TRANSLATE(t, trans)  (((t) & ~7) | ((trans) & 7))
#define SET_ENCODING(t, e)       (((t) & 15) | ((e) << 4))

```

The buffer-specific variable `translation-type` includes the current buffer's translation type as one of the above codes, combined with an encoding number that specifies the current buffer's encoding. Most functions for reading or writing a file take such a value as a `transl` parameter.

You can combine one of the above translation codes with an encoding number using the `MAKE_TRANSLATE()` macro. It takes a translation code `t`, an encoding code `e`, and a code `f` which, if nonzero, indicates that the resulting translation value was explicitly specified by the user in some way, not auto-detected.

Use the macros `GET_ENCODING()` and `GET_LINE_TRANSLATE()` to extract the encoding code or the line translation code, respectively, from a translation type value.

Use the macros `SET_ENCODING()` and `SET_TRANSLATE()` to combine an existing translation type value `t` with a new encoding code or line translation code, respectively.

```

user int default_translation_type;
user int new_buffer_translation_type;
int give_line_translate(char *fname)
int ask_line_translate()
get_fallback_translation_type()
buffer int fallback_translation_type;

```

A user can set the `default-translation-type` variable to one of the above translation codes to force Epsilon to use a specific translation when it reads an existing file. If this variable is set to its default value of `FILETYPE_AUTO`, Epsilon examines the file to determine a translation method. Setting this variable to any other value forces Epsilon to use that line translation method for all files. (The variable can specify only an encoding, or only a line translation, or both.)

When Epsilon creates a new buffer, it sets the buffer's `translation-type` variable to the value of the `new-buffer-translation-type` variable. Epsilon does the same when you try to read a file that doesn't exist. You can set this variable if you want Epsilon to examine existing files to determine their translation type, but create new files with a specific translation type. By default this variable is set to `FILETYPE_AUTO`, so the type for new buffers becomes `FILETYPE_UNIX` in Epsilon for Unix, and `FILETYPE_MSDOS` elsewhere.

The `give_line_translate()` subroutine defined in `files.e` helps to select the desired translation method and encoding. Many commands that read a user-specified file call it. If a numeric prefix argument was not specified, it returns the default translation type, often the value of the `default-translation-type` variable. (See that variable for details.) But if a numeric prefix argument was specified, it prompts the user for the desired translation type and encoding.

The `ask_line_translate()` subroutine is similar, but doesn't take a file name, and won't use any setting that might override the `default-translation-type` variable. New EEL code should use `give_line_translate()` instead.

The `get_fallback_translation_type()` primitive returns the translation type code Epsilon would assign when reading an empty file, or when writing a buffer whose translation type code was set to `FILETYPE_AUTO`. It returns `FILETYPE_UNIX` under Unix and `FILETYPE_MSDOS` on other platforms, but may be overridden for the current buffer by setting the buffer-specific `fallback_translation_type` variable.

10.3.4 Character Encoding Conversions

```
char *encoding_to_name(int enc)
int encoding_from_name(char *name)
```

When Epsilon reads or writes a file, it converts text between the Unicode character representation it uses internally and one of various file encodings. Epsilon represents each possible encoding with a number.

These numbers may change from one version of Epsilon to the next, so if an encoding setting must be recorded somehow, it should be recorded by name, not by number. Certain specific encodings will not change their codes: the encoding "auto-detect" is always numbered 0, and the encoding "raw" is always numbered 1.

The `encoding_from_name()` primitive returns the number of an encoding given its name. It returns -1 pointer if the encoding name is unknown.

The `encoding_to_name()` primitive returns the name of an encoding given its number. It returns a NULL pointer if the encoding number is unknown. Many encodings have more than one name, but this primitive treats each name as a separate encoding, even if it's an alias of another encoding.

```
int file_convert_write(char *file, int trans,
                      struct file_info *f_info)
int save_remote_file(char *fname, int trans,
                    struct file_info *finfo)
buffer char *(*file_io_converter)();
char *oem_file_converter(int func)
zeroed char *(*new_file_io_converter)();
zeroed buffer char file_write_newfile;
```

The `do_save_file()` subroutine uses the `file_convert_write()` subroutine to actually write the file. Like `new_file_write()`, it takes a file name, a line translation code as described under `translation-type`, and a structure which Epsilon will fill with information on the file's write date, file type, and so forth. See `do_save_file()` above for details.

Unlike primitives such as `new_file_write()`, the `file_convert_write()` subroutine knows how to handle URL files by calling the `save_remote_file()` subroutine.

In addition to the built-in conversion codes described above, Epsilon also supports user-defined EEL conversion routines. These are currently used only for DOS/OEM files read using the `find-oem-file` command and similar. The `file_convert_write()` subroutine handles writing these. It looks for a buffer-specific variable `file_io_converter`. This variable can be null, for no special translation, or it can contain a function pointer. For OEM files, for example, it points to the subroutine `oem_file_converter()`.

Any such subroutine will be called with a code indicating the desired action. The codes are defined in `eel.h`. The code `FILE_CONVERT_READ` tells the subroutine to translate the text in the current buffer as

appropriate when reading a file. The code `FILE_CONVERT_WRITE` tells the subroutine to translate the buffer as appropriate when writing a file.

Before actually performing a conversion, Epsilon will call the subroutine to ask if the conversion is safe (reversible), by passing the `FILE_CONVERT_ASK` in addition to one of the above flags. A conversion is reversible, and therefore safe, if the conversion followed by the opposite conversion (for instance, `ANSI => OEM => ANSI`) yields the original text. If the conversion isn't safe, the subroutine should ask the user for permission to proceed.

The converter should then return a null pointer to cancel the read or write operation, or any other value to let it proceed. You can add the `FILE_CONVERT_QUIET` flag, and the converter won't ask the user for confirmation, merely return a value indicating whether the conversion would be safe.

Whenever the `FILE_CONVERT_ASK` flag isn't present, the subroutine should return the name of its minor mode—Epsilon will display this in the mode line. The OEM converter returns " OEM".

When creating a new buffer, file-reading subroutines initialize the `file_io_converter` variable by copying the value of `new_file_io_converter`. Commands like `find-oem-file` temporarily set this variable to effect reading a file with OEM translation.

The `file_convert_write()` subroutine performs one more function. It checks the variable `file_write_newfile`. If this variable is nonzero, it arranges things so the attempt to write a file will fail with an error code if the file already exists, by passing the `FILE_IO_NEWFILE` code to `new_file_write()`.

```
int perform_unicode_conversion(int buf, int from, int to,
                             int flags, char *encoding)
```

The `perform_unicode_conversion()` primitive converts between 16-bit Unicode characters and various 8-bit encodings such as UTF-8. It converts characters in the range `from...to` in the specified buffer `buf` in place.

By default, the primitive converts from 16-bit Unicode characters to the named 8-bit encoding. The `CONV_TO_16` flag makes it convert in the opposite direction, from the specified 8-bit encoding to 16-bit characters.

The primitive returns the code `EBADENCODING` if it doesn't recognize the encoding name. It returns `ETOOBIG` when converting from 8-bit characters if one of the characters is outside the range 0–255. It returns 0 on success. The primitive moves point to the end of the buffer.

If the specified encoding has a defined signature (a byte order mark), and an entire buffer was converted, not just part of one, Epsilon adds the signature when converting to the encoding, and removes the signature, if there is one, when converting from the encoding.

```
int buffer_flags(int buf)
```

Internally, Epsilon stores the text of a buffer with 8 bits for each character, unless it contains some characters outside the range 0–255. In that case it uses 16 bits for each character. A buffer that once contained such characters but no longer does may still be stored as 16 bits per character. Epsilon transparently handles all needed translations between the two formats (for instance, when you copy text from one buffer to another), but it's occasionally useful to tell which format Epsilon is using.

The `buffer_flags()` primitive returns a bit mask. Check the bit represented by the `BF_UNICODE` macro; if it's present, the specified buffer `buf` is stored in 16-bit format internally. If `buf` is omitted or zero, the primitive checks the current buffer.

10.3.5 More File Primitives

```

user buffer short modified;
int get_buf_modified(int buf)
set_buf_modified(int buf, int val)
int unsaved_buffers()
zeroed buffer char save_without_prompt;
int is_unsaved_buffer()
int buffer_unchanged()
char *get_file_read_kibitz()

```

Epsilon maintains a variable that tells whether the buffer was modified since it was last saved to a file. The buffer-specific variable `modified` is set to 1 each time the current buffer is modified. It is set to 0 by the `file_read()`, `file_write()`, `new_file_read()`, and `new_file_write()` primitives, if they complete without error.

The `get_buf_modified()` and `set_buf_modified()` subroutines let you get or set the value of this variable for some other buffer `buf`.

The `unsaved_buffers()` subroutine defined in `files.e` returns 1 if there are any modified buffers. It doesn't count empty buffers, or those with no associated file names. If an EEL program creates a buffer that has an associated file name and is marked modified, but still doesn't require saving, it can set the buffer-specific variable `discardable_buffer` nonzero to indicate that the current buffer doesn't require any such warning. An EEL program can set the buffer-specific `save_without_prompt` variable nonzero for a buffer to have it silently saved without prompting, whenever Epsilon checks for unsaved buffers.

The `unsaved_buffers()` subroutine calls the `is_unsaved_buffer()` subroutine to check on an individual buffer. It tells if the current buffer shouldn't be deleted, and checks for the `discardable_buffer` variable as well as the `buffer-not-saveable` variable and other special kinds of buffers.

The `buffer_unchanged()` primitive returns a nonzero value if the current buffer has been modified since the last call of the `refresh()` or `maybe_refresh()` primitives. It returns zero if the buffer has not changed since that time. Epsilon calls `maybe_refresh()` to display the screen after each command.

The `get_file_read_kibitz()` primitive returns an explanatory message from the last time the current buffer was read, indicating why Epsilon chose a particular line translation. See the `file-read-kibitz` variable.

```

user buffer char *filename;
set_buffer_filename(char *file)
get_buffer_filename(char *fname)

```

The file reading and writing functions are normally used with the file name associated with each buffer, which is stored in the buffer-specific `filename` variable. To set this variable, use the syntax `filename = new value;`. Don't use `strcpy()`, for example, to modify it.

The `set_buffer_filename()` subroutine defined in `files.e` sets the file name associated with the current buffer. However, unlike simply setting the primitive variable `filename` to the desired value, this function also modifies the current buffer's name to match the new file name, takes care of making sure the file name is in absolute form, and updates the buffer's access "timestamp". The `bufed` command uses this timestamp to display buffers sorted by access time.

Some buffers such as `dired` buffers have an associated file name, but since they aren't copies of files, they don't store it in the `filename` variable. Call the `get_buffer_filename()` subroutine to copy the current buffer's associated file name to `fname`, whether it's stored in `filename` or not.

```

user int errno;
file_error(int code, char *file, char *unknown)
char no_popup_errors;

```

File primitives that fail often place an error code in the `errno` variable. The `file_error()` primitive takes an error code and a file name and displays to the user a textual version of the error message. It also takes a message to print if the error code is unknown.

Under MS-Windows, the `file_error()` primitive pops up a message box to report the error. If EEL code sets the `no_popup_errors` variable nonzero, Epsilon will display such messages in the echo area instead, as it does under other operating systems.

```

int do_insert_file(char *file, int transl) /* files.e */
int write_part(char *file, int transl, int start, int end)

```

The `do_insert_file()` subroutine inserts a file into the current buffer, like the `insert-file` command. The `write_part()` subroutine writes only part of the current buffer to a file. Each displays an error message if the file could not be read or written, and returns either an error code or 0.

```

locate_window(char *buf, char *file) /* buffer.e */
int buf_in_window(int bnum)
int is_buf_in_window(int bnum)
is_buffer_in_window(char *buf)
int count_windows_with_buf(int buf, int flags)

```

The `locate_window()` subroutine defined in `window.e` tries to display a given file or buffer by changing windows. If either of the arguments is an empty string "" it will be ignored. If a buffer with the specified name or a buffer displaying the specified file is shown in a window, the subroutine switches to that window. Otherwise, it makes the current window show the indicated buffer, if any.

The `buf_in_window()` primitive finds a window that displays a given buffer, and returns its window handle. It returns -1 if no window displays that buffer.

The `is_buf_in_window()` subroutine is similar, but excludes system windows. The `is_buffer_in_window()` subroutine takes a buffer name instead of a buffer number. They both return the handle of a non-system window displaying that buffer, or -1 if none.

The `count_windows_with_buf()` primitive returns the number of windows displaying the buffer. The flag bit 1 makes it skip system windows.

```

int delete_file(char *file)

```

The `delete_file()` primitive deletes a file. It returns 0 if the deletion succeeded, and -1 if it failed. The `errno` variable has a code describing the error in the latter case.

```

int rename_file(char *oldfile, char *newfile)

```

The `rename_file()` primitive changes a file's name. It returns zero if the file was successfully renamed, and nonzero otherwise. The `errno` variable has a code describing the error in the latter case. You can use this primitive to rename a file to a different directory, but you cannot use it to move a file to a different disk.

```

int copyfile(char *oldfile, char *newfile)

```

The `copyfile()` primitive makes a copy of the file named `oldfile`, giving it the name `newfile`, without reading the entire file into memory at once. The copy has the same time and date as the original. The primitive returns zero if it succeeds. If it fails to copy the file, it returns a nonzero value and sets `errno` to indicate the error.

```
int make_backup(char *file, char *backupname)
```

The `make_backup()` primitive does whatever is necessary to make a backup of a file. It takes the name of the original file and the name of the desired backup file, and returns 0 if the backup was made. Otherwise, it puts an error code in `errno` and returns a nonzero number. The primitive may simply rename the file, if this can be accomplished without losing any special attributes or permissions the original file has. If necessary, Epsilon copies the original file to its backup file.

```
int get_file_read_only(char *fname)
int set_file_read_only(char *fname, int val)
int set_file_opsys_attribute(char *fname, int attribute)
```

The `get_file_read_only()` primitive returns 1 if the file `fname` has been set read-only, 0 if it's writable, or -1 if the file's read-only status can't be determined (perhaps because the file doesn't exist). The `set_file_read_only()` primitive sets the file `fname` read-only (if `val` is nonzero) or writable (if `val` is zero). It returns 0 if an error occurred, otherwise nonzero.

Under Unix, `set_file_read_only()` sets the file writable for the current user, group and others, as modified by the current `umask` setting (as if you'd just created the file). Other permission bits aren't modified.

The `set_file_opsys_attribute()` primitive sets the raw attribute of a file. The precise meaning of the attribute depends on the operating system: under Unix this sets the file's permission bits, while in other environments it can set such attributes as Hidden or System. The primitive returns nonzero if it succeeds. See the `opsysattr` member of the structure set by `check_file()` to retrieve the raw attribute of a file.

```
int is_directory(char *str)
int url_is_a_directory(char *url)
int is_pattern(char *str)
```

The `is_directory()` primitive takes a string, and asks the operating system if a directory by that name exists. If so, `is_directory()` returns 1; otherwise, it returns 0. The `url_is_a_directory()` subroutine just calls `is_directory()` on local files, but also handles file names that start with `scp:` or `ftp:`. Since it may need to ask the remote system if a URL names a directory, it may not return quickly. Also see the `check_file()` primitive on page 505, and the `remote_file_type()` subroutine on page 512.

The `is_pattern()` primitive takes a string, and tells whether it forms a file pattern with wildcards that may match several files. It returns 2 if its file name argument contains the characters `*` or `?`. These characters are always wildcard characters and never part of a legal file name. The function returns 1 if its file name argument contains any of the following characters: left square-bracket, left curly-bracket, comma, or semicolon. These characters can sometimes be part of a valid file name (depending upon the operating system and file system in use), but are also used as file pattern characters in Epsilon. It returns 3 if the file name contains both types of characters, and it returns 0 if the file name contains none of these characters. (See the `file_match()` primitive to retrieve file names matching a pattern.)

```
user char file_pattern_wildcards;
#define FPAT_COMMA (1)
```

```

#define FPAT_SEMICOLON      (2)
#define FPAT_SQUARE_BRACKET (4)
#define FPAT_CURLY_BRACE    (8)
#define FPAT_ALL             (FPAT_COMMA | FPAT_SEMICOLON \
                               | FPAT_SQUARE_BRACKET | FPAT_CURLY_BRACE)

```

You can control which of the characters `[]{};` Epsilon will consider a wildcard character in file patterns by setting the `file-pattern-wildcards` variable. This affects the `do_dired()`, `is_pattern()`, `file_match()`, `dired_standardize()`, `check_file()`, and `is_directory()` primitives. Each bit in the variable enables a different set of characters.

`FPAT_COMMA` enables the `,` character, `FPAT_SEMICOLON` enables the `;` character, `FPAT_SQUARE_BRACKET` enables recognizing `[]` sequences, and `FPAT_CURLY_BRACE` lets Epsilon recognize `{}` sequences. The default value enables all these characters.

10.3.6 File Properties

```
int check_file(char *file, ?struct file_info *f_info)
```

The `check_file()` primitive gets miscellaneous information on a file or subdirectory from the operating system. It returns codes defined by macros in `codes.h`.

If its argument `file` denotes a pattern that may match multiple files, it returns `CHECK_PATTERN`. (Use the `file_match()` primitive described on page 571 to retrieve the matches. Note that `CHECK_PATTERN` checks only for the system-wide wildcard characters `*` and `?`, not for Epsilon's extended file pattern characters. See the `is_pattern()` primitive described on page 504 to check for those.)

If `file` names a directory or a file, it returns `CHECK_DIR` or `CHECK_FILE`, respectively. If `file` names a device, it returns `CHECK_DEVICE`. If `file` has the form of a URL, not a regular file, it returns `CHECK_URL`.

Under operating systems that support it, `check_file()` returns `CHECK_PIPE` for a named pipe and `CHECK_OTHER` for an unrecognized special file. Otherwise, it returns 0.

If `f_info` has a non-null value, `check_file()` fills the structure it points to with information on the file or directory, except when it returns 0 or `CHECK_URL`. The structure has the following format (defined in `eel.h`):

```

struct file_info {      /* used by check_file() */
    int fsize;           /* file size in bytes */
    int opsysattr;       /* system dependent attribute */
    int raw_file_date;   /* opsys-dependent date */
    short year;
    short month;         /* 1-12 */
    short day;           /* 1-31 */
    short hour;          /* 0-23 */
    short minute;        /* 0-59 */
    short second;        /* 0-59 */
    byte attr;           /* epsilon standardized attribute */
    byte check_type;     /* file/directory/device code */
};
#define ATTR_READONLY    1
#define ATTR_DIRECTORY   2

```

The `check_type` member contains the same value as `check_file()`'s return code. The `attr` member contains two flags: `ATTR_READONLY` if the file cannot be written, or `ATTR_DIRECTORY` if the operating system says the file is actually a directory. The `opsysattr` member contains a raw attribute code from the operating system: the meaning of bits here depends on the operating system, and Epsilon doesn't interpret them. (See the `set_file_opsys_attribute()` primitive to set raw attribute codes for a file.)

Epsilon also provides the timestamp of a file, in two formats. The interpreted format (year, month, etc.) uses local time, and is intended to match the file timestamp shown in a directory listing. By contrast, in most cases the raw timestamp won't be affected by a change in time zones, the arrival of daylight saving time, or similar things, as the interpreted format will be, so it's more convenient when comparing timestamps. Under some operating systems Epsilon doesn't provide a raw timestamp; this field will be zero in that case.

For the second parameter to `check_file()`, make sure you provide a *pointer* to a `struct file_info`, not the actual structure itself. You can omit this parameter entirely if you only want the function's return value.

```
unique_filename_identifier(char *fname, int id[3])
unique_file_ids_match(int a[3], int b[3])
```

The `unique_filename_identifier()` primitive takes a file name and fills the `id` array with a set of values that uniquely describe it. Two file names with the same array of values refer to the same file. (This can happen under Unix due to symbolic or hard links.) If the primitive sets `id[0]` to zero, no unique identifier was found; comparisons between two file names, one or both of which return `id[0]==0`, must assume that the names might or might not refer to the same file. At this writing only Epsilon for Unix supports this feature; in other versions, `unique_filename_identifier()` will always set `id[0]` to zero.

The `unique_file_ids_match()` subroutine compares two `id` arrays from `unique_filename_identifier()`, returning nonzero if they indicate the two file names supplied to `unique_filename_identifier()` refer to the same file, and zero if they do not, or Epsilon cannot determine this.

```
int compare_dates(struct file_info *a,
                  struct file_info *b)
format_date(char *msg, int year, int month,
             int day, int hour, int minute,
             int second)
format_file_date(char *s, struct file_info *p)
```

The `compare_dates()` subroutine defined in `filedate.e` can be used to compare the dates in two `file_info` structures. It returns 0 if they have the same date and time, nonzero if they differ.

The `format_date()` subroutine takes a date and converts it to text form, using the format specified by the `date-format` variable. The `format_file_date()` subroutine takes a `file_info` structure and converts it to text form by calling `format_date()`.

```
int check_dates(int save)                /* filedate.e */
```

The `check_dates()` subroutine defined in `filedate.e` compares a file's time and date on disk with the date saved when the file was last read or written. If the file on disk has a later date, it warns the user and asks what to do. Its parameter should be nonzero if Epsilon was about to save the file, otherwise zero. The function returns nonzero if the user said not to save the file.

The following example command uses `check_file()` to display the current file name and its date in the echo area.

```
#include "eel.h"

command show_file_date()
{
    struct file_info ts;

    if (check_file(filename, &ts))
        say("%s: %d/%d/%d", filename,
            ts.month, ts.day, ts.year);
}
```

10.3.7 Low-level File Primitives

```
int lowopen(char *file, int mode)
```

The following primitives provide low-level access to files. The `lowopen()` primitive takes the name of a file and a mode code. It returns a “file handle” for use with the other primitives. The mode may be 0 for reading only, 1 for writing only, or 2 for both. If the file doesn’t exist already, the primitive will return an error, unless you use mode 3. Mode 3 creates or empties the file first, and permits reading and writing.

```
int lowread(int handle, byte *buf, int count)
int lowwrite(int handle, byte *buf, int count)
```

The `lowread()` primitive tries to read the specified number of bytes, putting them in the byte array `buf`, and returns the number of bytes it was able to read. A value of 0 indicates the file has ended. The `lowwrite()` primitive tries to write the specified number of bytes from the byte array `buf`, and returns the number it was able to write. A return value different from `count` may indicate that the disk is full. See page 542 for functions to help translate between bytes and characters.

```
int lowseek(int handle, int offset, int mode)
int lowclose(int handle)
```

The `lowseek()` primitive repositions within the file. If the mode is 0, it positions to the *offset*th byte in the file, if 1 to the *offset*th byte from the previous position, and if 2 to the *offset*th byte from the end. The primitive returns the new offset within the file. Finally, the `lowclose()` primitive closes the file. All these routines return -1 if an error occurred and set `errno` with its code.

```
int lowaccess(char *fname, int mode)
#define LOWACC_R      4          /* file is readable. */
#define LOWACC_W      2          /* file is writable. */
#define LOWACC_X      1          /* file is executable. */
```

The `lowaccess()` primitive takes a file name and a code indicating whether the file’s read access, write access or execute access should be tested, or zero if only the file’s existence need be checked. It returns 0 if the file is accessible for the specified purpose (can be read, can be written, can be executed, exists), or -1 if not.

10.3.8 Directories

```
getcd(char *dir)
int chdir(char *dir)
```

The `getcd()` primitive returns the current directory, placing it in the provided string. For Windows, the format is `C:\harold\work`.

The `chdir()` primitive sets the current directory. (Under Windows, it sets the current drive as well if its argument refers to a drive. For example, invoking `chdir("A:\letters")`; sets the current drive to A, then sets the current directory for drive A to `\letters`. `chdir("A:")`; sets only the current drive.)

The result for this primitive is 0 if the attempt succeeded, and -1 if it failed. The `errno` variable is set with a code showing the type of error in the latter case.

```
put_directory(char *dir) /* files.e subr. */
int get_buffer_directory(char *dir)
```

The `put_directory()` subroutine copies the directory part of the file name associated with the current buffer into `dir`. Normally the directory name will end with a path separator character like `'/'` or `'\'`. If the current buffer has no associated file name, `dir` will be set to the empty string.

The `get_buffer_directory()` subroutine gets the default directory for the current buffer in `dir`. In most cases this is the directory part of the buffer's `filename` variable, but special buffers like `dired` buffers have their own rules. The subroutine returns nonzero if the buffer had an associated directory. If the buffer has no associated directory, the subroutine puts Epsilon's current directory in `dir` and returns 0.

```
user char *process_current_directory;
```

Epsilon stores the concurrent process's current directory in the `process_current_directory` variable. Setting this variable has no effect. Under Windows, Epsilon tries to detect the process's current directory from EEL code and set this variable. See the variable `use-process-current-directory` for more details. On other platforms, Epsilon sets up a modified shell prompt which sends a hidden current directory message, received internally.

```
int mkdir(char *dir)
int rmdir(char *dir)
```

The `mkdir()` subroutine makes a new directory with the given name, and the `rmdir()` subroutine removes an empty directory with the given name. Each primitive returns 0 on success and -1 on failure, and sets `errno` in the latter case, as with `chdir()`.

```
get_customization_directory(char *dir)
```

When Epsilon starts, it locates its customization directory, as described on page 13. The `get_customization_directory()` primitive copies the name of this directory to `dir`. The directory name always ends with a path separator character, either `'\'` or `'/'`.

Dired Subroutines

```

int dired_one(char *files)    /* dired.e */
int create_dired_listing(char *files)
int make_dired(char *files)
int do_remote_dired(char *files)
int do_dired(char *files)
int is_dired_buf()          /* dired.e */

```

The `dired_one()` subroutine takes a file name pattern as its argument and acts just like the `dired` command does, making a dired buffer, filling it and putting it in dired mode. It puts its pattern in a standard form and chooses a suitable buffer name, then calls the `create_dired_listing()` subroutine. This function prepares the buffer and displays suitable messages, then calls `make_dired()`.

The `make_dired()` subroutine handles FTP dired requests by calling `do_remote_dired()`, and passes local dired requests to the `do_dired()` primitive to fill the buffer with directory information.

Each of these routines takes a file name with wildcard characters such as `*` and `?`, and inserts in the current buffer exactly what the `dired` command does (see page 139). Each returns 0 normally, and 1 if there were no matches.

By default, the `do_dired()` primitive ignores the abort key. To permit aborting a long file match, set the primitive variable `abort_file_matching` using `save_var` to tell Epsilon what to do when the user presses the abort key. See page 572 for details.

The `is_dired_buf()` subroutine returns 1 if the current buffer is a dired buffer, otherwise 0.

```

dired_standardize(char *files)
standardize_remote_pathname(char *files)
remote_dirname_absolute(char *dir)
drop_dots(char *path)

```

Sometimes there are several interchangeable ways to write a particular file pattern. For example, `/dir1/dir2/*` always makes the same list of files as `/dir1/dir2/` or `/dir1/dir2`. The `dired_standardize()` primitive converts a dired pattern to its simplest form, in place. In the example, the last pattern is considered the simplest form.

The `standardize_remote_pathname()` subroutine is similar, but operates on FTP and SCP URLs. It calls several other subroutines to help.

The `remote_dirname_absolute()` subroutine converts a relative remote pathname to an absolute one in place. It performs an FTP or SCP operation to get the user's home directory, then inserts it into the given pathname.

The `drop_dots()` subroutine removes `.` and interprets `..` in a pathname, modifying it in place. It removes any `..` components at the start of a path.

```

detect_dired_format()
zeroed buffer char dired_format;
#define DF_UNIX      1
#define DF_SIMPLE    2
#define DF_OLDNT     3
#define DF_VMS       4
int get_dired_item(char *prefix, int func)

```

The `direcd` command supports several different formats for directory listings. Besides the standard format it uses for local directory listings, as generated by the `do_dired()` primitive, it understands the directory listings generated by FTP servers that run on Unix systems (and the many servers on other operating systems that use the same format), as well as several others.

The `detect_dired_format()` subroutine determines the proper format by scanning a `direcd` buffer, and sets the `direcd_format` variable as appropriate. A value of 0 indicates the default, local directory format. The other values represent other formats.

Various subroutines in `direcd` use the `get_dired_item()` subroutine to help locate format-specific functions or variables, to do tasks that depend on the particular format. The subroutine takes a prefix like “`direcd-isdir-`” and looks for a function named `direcd-isdir_unix()` (assuming the `direcd_format` variable indicates Unix). It returns the name table index of the function it found, if there is one, or zero otherwise.

If its parameter `func` is nonzero, it looks only for functions; if zero, it looks only for variables. You can use an expression like `(* (int (*()) i) ())` to call the function (assuming `i` is the value returned by `get_dired_item()`), or an expression like `get_str_var(i)` to get the value of a variable given its index.

10.3.9 Manipulating File Names

```
absolute(char *file, ?char *dir)
relative(char *abs, char *rel, ?char *dir)
int is_relative(char *fname)
```

Because the current directory can change, through use of the `chdir()` primitive described above, Epsilon normally keeps file names in absolute pathname form, with all the defaults in the name made explicit. It converts a file name to the appropriate relative pathname whenever it displays the name (for example, in the mode line).

The `absolute()` primitive takes a pointer to a character array containing a file name. It makes the file name be an absolute pathname, with all the defaults made explicit. For example, if the default drive is B:, the current directory is `/harold/papers`, the `path_sep` variable is ‘`\`’ and the 80 character array `fname` contains “`proposal`”; calling `absolute()` with the argument `fname` makes `fname` contain “`B:\harold\papers\proposal`”.

The primitive `relative()` does the reverse. It takes a file name in absolute form and puts an equivalent relative file name in a character array. Unlike `absolute()`, which modifies its argument in place, `relative()` makes a copy of the argument with the changes. If the default drive is B:, the current directory is `\harold` and the 80 character array `abs` contains `B:\harold\papers\proposal`, calling `relative(abs, rel)`; puts “`papers\proposal`” in the string array `rel`. You can also get a relative file name by using the `%r` format specifier in any Epsilon primitive that accepts a printf-style format string.

The `relative()` and `absolute()` primitives each take an optional additional argument, which names a directory. The `absolute()` primitive assumes that any relative file names in its first argument are relative to the directory named by the second argument. (If the second argument is missing or null, the primitive assumes that relative file names are relative to the current directory.) Similarly, if you provide a third argument to the `relative()` primitive, it makes file names relative to the specified directory, instead of the current directory.

Note that in EEL string or character constants, the `\` character begins an escape sequence, and you must double it if the character `\` is to appear in a string. Thus the Windows file name `\harold\papers` must appear in an EEL program as the string “`\\harold\\papers`”.

The `is_relative()` primitive returns nonzero if the file name looks like a relative pathname, not an absolute pathname. (It’s not intended for use with URLs.)

```
char *get_tail(char *file, ?int dirok)
```

The `get_tail()` primitive takes a string containing a file name and returns a pointer to a position in the string after the name of the last directory. For example, suppose that `file` is the string `“/harold/papers/proposal”`. Then

```
get_tail(file, 0)
```

would return a pointer to `“proposal”`. Since the pointer returned is to the original string, you can use this primitive to modify that string. Using the above example, a subsequent

```
strcpy(get_tail(file, 0), "sample");
```

would make `file` contain the string `“/harold/papers/sample”`. The `dirok` argument says what to do with a file name ending with a separator character `‘\’` or `‘/’`. If `dirok` is nonzero the primitive returns a pointer to right after the final separator character. If `dirok` is zero, however, the primitive returns a pointer to the first character of the final directory name. (If `file` contains no directory name, the primitive returns a pointer to its first character when `dirok` is zero.)

```
char *get_extension(char *file)
```

The `get_extension()` primitive returns a pointer to the final extension of the file name given as its argument. For example, an invocation of

```
get_extension("text.c")
```

would return a pointer to the `“.c”` part, and `get_extension("text")` would return a pointer to the null character at the end of the string. Like `get_tail()`, you can use this primitive to modify the string.

```
int is_path_separator(int ch)
```

The `is_path_separator()` primitive tells if a character is one of the characters that separate directory or drive names in a file name. It returns 1 if the character is `‘\’` or `‘/’`, 2 if the character is `‘.’`, otherwise 0. Under Unix, it returns 1 if the character is `‘/’`, otherwise 0.

```
user char path_sep;
```

The `path_sep` variable contains the character for separating directory names. It is `‘\’` under Windows, `‘/’` under Unix.

```
add_final_slash(char *fname)
drop_final_slash(char *fname)
```

The `add_final_slash()` primitive adds a path separator character like `/` or `\` to the end of `fname`, if there isn't one already. The `drop_final_slash()` primitive removes the last character of `fname` if it's a path separator. These primitives never count `:` as a path separator.

```
abbreviate_file_name(char *file, int room)
```

The `abbreviate_file_name()` subroutine defined in `disp.e` modifies the filename `file` so it's no more than `room` characters long, by replacing sections of it with an ellipsis (`...`). If `file` is no more than `room` characters long to begin with, it won't be changed. Values of `room` less than 10 will be treated as 10.

```
int is_remote_file(char *fname)
char url_services[50] = "ftp|http|telnet|scp|ssh";
int remote_file_type(char *fname)
```

The `is_remote_file()` primitive tells whether `fname` looks like a valid URL. It returns 1 if `fname` starts with a service name like `ftp://`, `http://`, or `telnet://`, or 2 if `fname` appears to be an Emacs-style remote file name like `/hostname:filename`. It uses the `url_services` variable to determine which service names are valid; this must be a series of |-separated names.

The `remote_file_type()` subroutine is somewhat similar; it tries to determine if a file `fname` refers to a remote directory, a file pattern, or some other sort of thing. It returns 1 if `fname` doesn't have the format of a remote file (so it might be a local file), 2 if its syntax is invalid, 3 if it's a remote file that specifies a service other than `ftp` or `scp`, 4 if there's no file name after its host name, 5 if it uses wildcards, 6 if it ends in a path separator, or 7 if it uses `~` to name a user's home directory and has no file name following that.

If none of these cases apply, the subroutine contacts the remote system to test whether `fname` refers to a directory or a file, and returns 8 if it's a directory, otherwise 0. (A value of 0 doesn't indicate there's necessarily a file by that name, just that there is no directory by that name.) While Epsilon is first starting up, the subroutine returns 9 instead of contacting the remote system, since Epsilon can't respond to a site's password request until it has finished initializing, and this synchronous test could otherwise block it from finishing its initialization..

```
get_executable_directory(char *dir)
get_executable_file(char *dest, char *prog, int quoted)
```

The `get_executable_directory()` function stores the full pathname of the directory containing the Epsilon executable into `dir`. The `get_executable_file()` function uses this; it writes into `dest` the full pathname of a file `prog` in the same directory as Epsilon's executable. If `quoted` is nonzero, the file name is inside a pair of quote characters, for use in a command line.

```
look_up_tree(char *res, char *file, char *dir, char *stop)
int is_in_tree(char *file, char *tree) /* files.e subr. */
```

The `look_up_tree()` subroutine searches for `file` in the given directory `dir`, its parent directory, and so forth, until it finds a file named `file` or reaches the root directory. If it finds such a file, it returns nonzero and puts the absolute pathname of the file into the character array `res`. If it doesn't find a file with the given name, it returns zero and leaves `res` set to the last file it looked for. If `file` is an absolute pathname to begin with, it puts the same file name in `res`, and returns nonzero if that file exists. If `dir` is a null pointer, `look_up_tree()` begins at the current directory. If `stop` is non-null, the function only examines child directories of the directory `stop`. The function stops as soon as it reaches a directory other than `stop` or one of its subdirectories. This function assumes that all its parameters are in absolute pathname form.

The `is_in_tree()` subroutine returns nonzero if the pathname `file` is in the directory specified by `dir` or one of its subdirectories. Both of its parameters must be in absolute pathname form.

```
user char path_list_char;
```

The `path_list_char` variable contains the character separating the directory names in a configuration variable like `EPSPATH`. It is normally `';`, except under Unix, where it is `'.'`.

```
build_filename(char *result, char *pattern, char *file)
```

The `build_filename()` subroutine constructs file names from name templates (see page 122). It copies pattern to result, replacing the various % template codes with parts of file, which it obtains by calling primitives such as `get_tail()` and `get_extension()`. The `expand_string_template()` subroutine on page 541 provides a more generalized facility to do this.

```
int fnamecmp(char *f1, char *f2)      /* buffer.e */
int filename_rules(char *fname)
```

The `fnamecmp()` subroutine compares two file names like the `strcmp()` primitive, returning 0 if they're equal, a positive number if the first comes before the second, or a negative number otherwise. However, it does case-folding on the file names first if this is appropriate for the particular file systems.

The `filename_rules()` primitive asks the operating system if a certain file system is case-sensitive or case-preserving, and returns other information too. It takes the name of any file or directory (which doesn't have to exist) on the file system, and returns a code whose values are represented by macros defined in `codes.h`. See page 126 for more information on how Epsilon determines the appropriate code for each file system.

The `FSYS_CASE_IGNORED` code indicates a non-case-preserving file system like DOS. The `FSYS_CASE_PRESERVED` code indicates a case-preserving file system like NTFS or VFAT. The `FSYS_CASE_SENSITIVE` code indicates a case-sensitive file system like Unix. The `FSYS_CASE_UNKNOWN` code indicates that Epsilon couldn't determine anything about the file system.

The function also returns a bit flag `FSYS_SHORT_NAMES`, valid whenever any code but `FSYS_CASE_UNKNOWN` is returned, that indicates whether only 8+3 names are supported. Use the mask macro `FSYS_CASE_MASK` to strip off this bit: for example, the expression

```
(filename_rules(f) & FSYS_CASE_MASK) == FSYS_CASE_SENSITIVE
```

is nonzero if the file system is case-sensitive.

The primitive also may return a bit indicating the type of drive a file is located on, if Epsilon can determine this. `FSYS_NETWORK` indicates the file is on a different computer and is being accessed over a network. `FSYS_CDROM` indicates the file is on a CD-ROM disk. `FSYS_REMOVABLE` indicates the file is on a removable medium like a floppy disk or Zip disk. And `FSYS_LOCAL` indicates the file is on a local (non-network) hard disk. At most one of the these bits will be present.

Epsilon for Unix returns `FSYS_CASE_SENSITIVE` for all files, even if they happen to lie on a file system that might use different rules natively. It can't detect the type of drive a file is on either.

```
int ok_file_match(char *s)           /* complete.e */
```

The `ok_file_match()` subroutine checks a file name to see if the `ignore_file_extensions` variable should exclude it from completion. It returns 0 if the file name should be excluded, or 1 if the file name is acceptable.

```
char *lookpath(char *file, ?int curdir)
char *look_on_path(char *file, int flags, char *path, ?int skip)
```

The `lookpath()` primitive looks in various standard Epsilon directories for a readable file with the supplied name. As soon as Epsilon locates the file, it returns the file's name. If it can't find the file, it returns a null pointer. See page 12 for more information on Epsilon's searching rules. The `look_on_path()` primitive is similar, but you can specify the path to use, and it offers some additional flexibility. These primitives will be described together.

First (for either primitive), if the specified file name is an absolute pathname, Epsilon simply checks to see if the file exists, and returns its name if it does, or a null pointer otherwise.

Next, if you call `lookpath()` with its optional parameter `curdir` nonzero (or if you call `look_on_path()` with the flag `PATH_ADD_CUR_DIR`), Epsilon looks for the file in the current directory. If `curdir` is zero or omitted (or `PATH_ADD_CUR_DIR` isn't specified), Epsilon skips this step (unless the file name explicitly refers to the current directory, like `".\filename"`).

The `lookpath()` primitive next looks for the file as explained on page 12, looking along your `EPSPATH`, or a default one.

Similarly, `look_on_path()` searches the provided path, which is in the same format as an `EPSPATH`, a list of directory names separated by semicolons for Windows, colons for Unix. The `PATH_ADD_EXE_DIR` bit makes it search in the executable's directory, like `-w32` does for `lookpath()`. The `PATH_ADD_EXE_PARENT` bit makes it search the executable's parent directory. It does both of these additional checks, when enabled, in the above order and just before searching the given path.

By default, `look_on_path()` only searches for files with the specified name. Add the `PATH_PERMIT_DIRS` flag if you want it to also return directories with that name. With the `PATH_PERMIT_WILDCARDS` flag, you can use a file pattern like `*.c` as the file name. The primitive will return the first matching file name.

If you supply `look_on_path()` with an optional `skip` parameter of *n*, it will skip over the first *n* matches it finds (so long as its parameter is a relative pathname). You can use this to reject a file and look for the next one on a path.

The value returned by each of these functions is only valid until the next time you call one of them. Copy the returned file name if you want to preserve it.

```
convert_to_8_3_filename(char *fname, ?int from8_3)
```

Under Windows, the `convert_to_8_3_filename()` primitive modifies the given file name by converting all long file names in `fname` to their short "8.3" file name aliases. Each component of a short file name has no more than eight characters, a dot, and no more than three more characters. For example, the file name `"c:\Windows\Start Menu\Programs\Windows Explorer.lnk"` might be translated to an equivalent file name of `"c:\Windows\STARTM~1\Programs\WINDOW~1.LNK"`. If the optional `from8_3` argument is nonzero, Epsilon translates in the reverse direction. Non-Windows versions of Epsilon will not modify the file name.

10.3.10 Internet Primitives

```
int telnet_host(char *host, int port, char *buf)
telnet_send(int id, char *text)
do_telnet(char *host, int port, char *buf)
buffer int telnet_id;
int telnet_server_echoes(int id)
```

Epsilon provides various commands that use Internet FTP, Telnet and similar protocols. This section documents how some parts of this interface work.

All Internet-related functions that retrieve text in the background and insert it in a buffer do so at the buffer's `type_point`, just like Epsilon's process functions.

First, Epsilon provides the primitives `telnet_host()` and `telnet_send()` for use with the Telnet protocol. The `telnet_host()` function establishes a connection to a host on the specified port, and using

the indicated buffer. It returns an identification code. The `telnet_send()` function can use this code to send text to the host. To kill the telnet job, call `telnet_send()` and pass `NULL` as the text. Commands normally call the `telnet_host()` function through the `do_telnet()` subroutine, which records the telnet identification code in the buffer-specific `telnet_id` variable, and does other housekeeping tasks.

The `telnet_server_echoes()` primitive accepts a telnet identification code as above, and returns 1 if the server on that connection is currently set to echo characters sent to it, or 0 if it is not.

```
int finger_user(char *user, char *host, char *buf)
int http_retrieve(char *resource, char *host, int port,
                 char *auth, char *buf, int flags)
char *http_force_headers;
int download_file_to_disk(char *url, char *fname, int sz)
show_url(char *url)
try_show_url(char *url)
```

The `finger_user()` primitive uses the Finger protocol to retrieve information on a particular user (if the host is running a Finger server). It takes the user name, the host, and the name of a buffer in which to put the results.

The `http_retrieve()` primitive uses the HTTP protocol to retrieve a page from a web site. It takes a resource name (the final part of a URL), a host, port, an authorization string (for password-protected pages) and destination buffer name, plus a set of flags. The `HTTP_RETRIEVE_WAIT` flag tells the function not to return until the transfer is complete. Without this flag the function begins the transfer and lets it continue in the background. The `HTTP_RETRIEVE_ONLY_HEADER` flag tells the function to retrieve only the header of the web page, not the body. Without this flag Epsilon will retrieve both; the first blank line retrieved separates the two.

If the `http_force_headers` variable is non-null and non-empty, `http_retrieve()` uses its contents for the HTTP request it sends to the remote system. It should contain a complete, valid HTTP request. But if it starts with a `+` character, then Epsilon simply adds the rest of its contents to each HTTP request. It should contain a series of header lines, each terminated by `\r\n`.

The `download_file_to_disk()` subroutine retrieves the document from the specified `url`, then writes it to the disk file named `fname`. It returns 0 on success, 1 if retrieving the file failed, or an error code from writing the file. (A 404 or other numeric error when retrieving a web page is treated as a success, as long as it returns a page of some sort. Check the HTTP Headers buffer for the error code if needed.) The `sz` parameter should be the file's expected size; if nonzero, the subroutine displays progress messages.

The `show_url()` subroutine displays the specified URL in a web browser, aborting with an error if it couldn't. The similar `try_show_url()` subroutine displays the URL in a browser, returning zero if it couldn't, nonzero if it could. A success code indicates merely that Epsilon was able to find a browser, not that the specified page exists.

```
int is_remote_buffer(int buf)
buffer char *buffer_url;
```

The `is_remote_buffer()` subroutine returns nonzero if the specified buffer has a telnet or ssh session running, or whose file name refers to a remote file (one accessed via scp or ftp).

In buffers with a telnet or ssh session, Epsilon sets the buffer-specific `buffer_url` variable to the URL used to create it. This is so it can restart the session later if necessary.

```
int ftp_op(char *buf, char *log, char *host, int port,
           char *usr, char *pwd, char *file, int op)
```

```
int do_ftp_op(char *buf, char *host, char *port,
             char *usr, char *pwd, char *file, int op)
```

The `ftp_op()` primitive uses the FTP protocol to send or retrieve files or get directory listings. It takes the destination or source buffer name, the name of a log buffer, a host computer name and port number, a user name and password, a file name, and an operation code that indicates what function it should perform (see below).

The `do_ftp_op()` subroutine is similar to `ftp_op()`, but it chooses the name of an appropriate FTP Log buffer, instead of taking the name of one as a parameter. Also, it arranges for the appropriate `ftp_activity()` function (see below) to be called, arranges for character-coloring the log buffer, and initializes the `ftp_job` structure that Epsilon uses to keep track of each FTP job.

The `FTP_RECV` operation code retrieves the specified file and the `FTP_SEND` code writes the buffer to the specified file name. The `FTP_LIST` code retrieves a file listing from the host of files matching the specified file pattern or directory name. The `FTP_MISC` code indicates that the file name actually contains a series of raw FTP commands to execute after connecting and logging in, separated by newline characters. Epsilon will execute the commands one at a time.

You can combine one of the above codes with some bit flags that modify the operation. Use the `FTP_OP_MASK` macro to mask off the bit flags below and extract one of the operation codes above.

Normally `ftp_op()` returns immediately, and each of these operations is carried out in the background. Add the code `FTP_WAIT` to any of the above codes, and the subroutine will not return until the operation completes.

The `FTP_ASCII` bit flag modifies the `FTP_RECV` and `FTP_SEND` operations. It tells Epsilon to perform the transfer in ASCII mode. By default, all FTP operations use binary mode, and Epsilon performs any needed line translation itself. But this doesn't work on some host systems (VMS systems, for example). See the `ftp-ascii-transfers` variable for more information.

The `FTP_USE_CWD` bit flag modifies how Epsilon uses the file name provided for operations like `FTP_RECV`, `FTP_SEND`, and `FTP_LIST`. By default, Epsilon sends the file name to the host as-is. For example, if you try to read a file `dirname/another/myfile`, Epsilon sends an FTP command like `RETR dirname/another/myfile`. Some hosts (such as VMS) use a different format for directory names than Epsilon's dired directory editor understands. So with this flag, Epsilon breaks a file name apart, and translates a request to read a file such as `dirname/another/myfile` into a series of commands to change directories to `dirname`, then to `another`, and then to retrieve the file `myfile`. The `FTP_PLAIN_LIST` bit flag makes `FTP_LIST` operations send a `LIST` command; without it, they send a `LIST -a` command so the remote system includes hidden files. The `ftp-compatible-dirs` variable controls these bits.

```
int url_operation(char *file, int op)
```

The `url_operation()` subroutine parses a URL and begins an Internet operation with it. It takes the URL and an operation code as described above for `ftp_op()`. If the code is `FTP_RECV`, then the URL may indicate a service type of `telnet://`, `http://`, or `ftp://`, but if the code is `FTP_SEND` or `FTP_LIST`, the service type must be `ftp://`. It can modify the passed URL in place to put it in a standard form. It calls one of the functions `do_ftp_op()`, `http_retrieve()`, or `do_telnet()` to do its work. It returns 0 on success. On failure it either returns nonzero or aborts, depending on the type of error.

```
ftp_misc_operation(char *url, char *cmd)
```

The `ftp_misc_operation()` subroutine uses the `do_ftp_op()` subroutine to perform a series of raw FTP commands. It takes an `ftp://` URL (ignoring the file name part of it) connects to the host, logs in, and then executes each of the newline-separated FTP commands in `cmd`. Dired uses this function to delete or move a group of files.


```
buffer int (*when_net_activity)();
net_activity(int activity, int buf, int from, int to)
```

As Epsilon performs Internet functions, it calls an EEL function to advise it of its progress. The buffer-specific variable `when_net_activity` contains a function pointer to the function to call. Epsilon uses the value of this variable in the destination buffer (or, in the case of the `NET_LOG_WRITE` and `NET_LOG_DONE` codes below, the log buffer). If the variable is zero in a buffer, Epsilon won't call any EEL function as it proceeds.

The EEL function will always be called from within a call to `getkey()` or `delay()`, so it must save any state information it needs to change, such as the current buffer, the position of point, and so forth, using `save_var`. The subroutine `net_activity()` shown above indicates what parameters the function should take—there's not actually a function by that name.

The `activity` parameter indicates the event that just occurred. A value of `NET_RECV` indicates that Epsilon has just received some characters and inserted them in a buffer. The `buf` parameter tells which buffer is involved. The `from` and `to` values indicate the new characters. A value of `NET_DONE` means that the net job running in buffer `buf` has finished. The above are the only activity codes generated for HTTP, Telnet, or Finger jobs.

FTP jobs have some more possible codes. `NET_SEND` indicates that another block of text has been sent. In this case, `from` indicates that number of bytes sent already from buffer `buf`, and `to` indicates the total number of bytes to be sent. The code `NET_LOG_WRITE` indicates that some more text has been written to the log buffer `buf`, in the range `from...to`. Finally, the code `NET_LOG_DONE` indicates that the FTP operation has finished writing to the log buffer. It occurs right after a `NET_DONE` call on FTP jobs.

```
ftp_activity(int activity, int buf, int from, int to)
finger_activity(int activity, int buf, int from, int to)
telnet_activity(int activity, int buf, int from, int to)
buffer int (*buffer_ftp_activity)();
```

The file `epsnet.e` defines the `when_net_activity` functions shown above, which provide status messages and similar things for each type of job. The `ftp_activity()` subroutine also calls a subroutine itself, defined just like these functions, through the buffer-specific variable `buffer_ftp_activity`. The `direcd` command uses this to arrange for normal FTP activity processing when retrieving directory listings, but also some processing unique to `direcd`.

```
int gethostname(char *host, ?int method)
```

The `gethostname()` primitive sets `host` to the computer's host name and returns 0. If it can't for any reason, it returns 1 and sets `host` to "?".

The `method` parameter controls which type of host name Epsilon retrieves. By default, it uses the machine's locally-configured host name. A value of 3 makes it instead retrieve the host's fully qualified domain name using DNS. Values of 1 or 2 make it do this only under Windows or Unix, respectively. Retrieving the DNS name in some network configurations can cause on-demand auto-dialing or delays if the machine's DNS server isn't accessible.

Parsing URLs

```
int prepare_url_operation(char *file, int op, struct url_parts *parts)
int get_password(char *res, char *host, char *usr)
int parse_url(char *url, struct url_parts *p)
int divide_url(char *url, struct url_parts *p)
```

Several subroutines handle parsing URLs into their component parts. These parts are stored in a `url_parts` structure, which has fields for a URL's service (`http`, `ftp`, and so forth), host name, port, user name if any, password if any, and the "file name": the final part of a URL, that may be a file name, a web page name or something else. Since an empty user name or password is legal, but is different from an omitted one, there are also fields to specify if each of these is present.

The `prepare_url_operation()` subroutine parses a URL and fills one of these structures. It aborts with an error if it doesn't recognize the service name, or if the service is something other than FTP but the operation isn't reading. The operation code is one of those used with the `ftp_op()` subroutine described on page 516. For example, it complains if you try to perform an `FTP_LIST` operation with a `telnet://` URL. It also prompts for a password if necessary, and saves the password for later use, by calling the `get_password()` subroutine. The subroutine returns zero on success, nonzero on failure.

The `get_password()` subroutine gets the password for a particular user/host combination. Specify the user and host, and the subroutine will fill in the provided character array `res` with the password. The first time it will prompt the user for the information; it will then store the information and return it without prompting in future requests. The subroutine is careful to make sure the password never appears in a state file or session file. To discard a particular remembered password, pass `NULL` as the first parameter. The next time `get_password()` is asked for the password of that user on that host, it will prompt the user again. The subroutine returns 1 if it was unable to prompt for a password because Epsilon was already prompting for some other text, and zero on success.

The `prepare_url_operation()` subroutine calls the `parse_url()` subroutine to actually parse the URL into a `url_parts` structure. The latter returns zero if the URL is invalid, or nonzero if it appears to be legal. (For parsing purposes, a partial URL like `http://` is considered legal, but with an empty host name.)

The `divide_url()` subroutine is similar to `parse_url()`, but doesn't divide the host section into its component parts. Like `parse_url()`, it returns zero if the URL is invalid, or nonzero if it appears to be legal.

For example, given the URL `scp://bob:secret%2Fcode@example.com:1022/path/to/file`, both `divide_url()` and `parse_url()` set the `service` member of the `url_parts` structure to "`scp`" and the `fname` member to "`path/to/file`".

But `divide_url()` then sets the `host` member to "`bob:secret%2Fcode@example.com:1022`", whereas `parse_url()` sets `host` to "`example.com`", `port` to "`1022`", `usr` to "`bob`", and `pwd` to "`secret/code`", also setting the `have_password` and `have_usr` members nonzero since the URL specified both. Notice that `parse_url()` decodes any %-escaped sequences in the user name or password sections, changing `%2F` to `/` in this example.

```
int split_string(char *part1, char *cs, char *part2)
int reverse_split_string(char *part1, char *cs, char *part2)
```

The `parse_url()` subroutine uses two helper subroutines. The `split_string()` subroutine divides a string `part1` into two parts, by searching it for one of a set of delimiter characters `cs`. It finds the first character in `part1` that appears in `cs`. Then it copies the remainder of `part1` to `part2`, and removes the delimiter character and the remainder from `part1`. It returns the delimiter character it found. If no delimiter character appears in `part1`, it sets `part2` to "" and returns 0. The `reverse_split_string()` subroutine is almost identical; it just searches through `part1` from the other end, and splits the string at the last character in `part1` that appears in `cs`.

```
char *get_url_file_part(char *url, int sep)
```

The `get_url_file_part()` subroutine helps to parse URLs. It takes a URL and returns a pointer to a position within it where its file part begins. For example, in the URL

<http://www.lugaru.com/why-lugaru.html>, the subroutine returns a pointer to the start of “why”. If `sep` is nonzero, the subroutine instead returns a pointer to the / just before “why”. If its parameter is not a URL, the subroutine returns a pointer to its first character.

10.3.11 Tagging Internals

This section describes how to add tagging support to Epsilon for other languages. Epsilon already knows how to find tags in C and EEL files, and in assembly languages files.

```
tag_suffix_ext()      /* example function */
tag_suffix_none()
tag_suffix_default()
tag_mode_c()
```

When Epsilon wants to add tags for a file, it first looks at the file’s extension and constructs a function name of the form `tag_suffix_ext()`, where *ext* is the extension. It tries to call this function to tag the file. If the file has no extension, it tries to call `tag_suffix_none()`.

If there is no function with the appropriate name, Epsilon looks for a function based on the current buffer’s mode. It constructs a function name of the form `tag_mode_mode()`, where *mode* is the value of the `major_mode` variable in the current buffer. If there is no mode-based function either, Epsilon calls `tag_suffix_default()` instead.

Thus, to add tagging for a language that uses file names ending in `.xyz`, define a function named `tag_suffix_xyz()`. Or if such files (and perhaps files with other extensions) use mode named “Xyz”, define a function named `tag_mode_xyz()`. In most cases, a mode-based name is more convenient.

```
add_tag(char *func, int pos)
```

The tagging function will be called with point positioned at the start of the buffer to be tagged. (Epsilon preserves the old value of point.) It should search through the buffer, looking for names it wishes to tag. To add a tag, it should call the subroutine `add_tag()`, passing it the tag name and the offset of the first character of the name within the file. You can use the tagging functions for C and assembler as examples to write your own tagging functions. They are in the source file `tags.e`.

The `pluck-tag` command uses a regular expression pattern to parse an identifier in the buffer. By default, it uses the pattern in the variable `tag-pattern-default`. A mode can define a variable like `tag-pattern-perl` or `tag-pattern-c` to make Epsilon use a different pattern. (For instance, the pattern for C mode says that identifiers can include `::` to specify a class name.)

Epsilon constructs a variable name, like `tag-pattern-perl`, from the current mode’s name. If a variable by that name exists, `pluck-tag` uses it in place of `tag-pattern-default`.

10.4 Operating System Primitives

10.4.1 System Primitives

```
char *getenv(char *name)
get_substitution_variables(int flags)
putenv(char *name)
char *verenv(char *name)
```

Use the `getenv()` primitive to return entries from the environment. The primitive returns a null pointer if no environment variable name exists. For example, after the Windows command “set waldo=abcdef”, the expression `getenv("waldo")` will return the string “abcdef”.

The `get_substitution_variables()` primitive inserts all environment variables and their values into the current buffer, much like the “set” command in Windows or the “env” command on other platforms, whenever the 1 bit is specified in its *flags* parameter. The 2 bit makes it include Windows shell folder names and their values.

The `putenv()` primitive puts strings in the environment. Normally environment entries have the form “NAME=definition”. This primitive manipulates Epsilon’s copy of the environment, which is passed on to any program that Epsilon runs, but it doesn’t affect the environment you get when you exit from Epsilon. The value of the argument to `putenv()` is evaluated later, when you actually invoke some other program from within Epsilon. For this reason, it is important that the argument to `putenv()` not be a local variable. Use `putenv(strkeep(value))`; to conveniently preserve the setting.

The `verenv()` primitive gets configuration variables. Epsilon for Windows looks in the system registry. Under Unix it retrieves the variables from the environment, like `getenv()`.

Regardless of the operating system, this primitive looks for alternate, version-specific forms of the specified configuration variable. For example, in version 7.0 of Epsilon, `verenv("MYVAR")` would return the value of a variable named MYVAR70, if one existed. If not, it would try the name MYVAR7. If neither existed, it would return the value of MYVAR (or a null pointer if none of these variables were found). See page 10 for complete information on configuration variables.

```
short opsys;
#define OS_WINDOWS 1    /* All Windows versions */
#define OS_UNIX    3    /* Includes Linux, FreeBSD, Mac OS. */
```

The `opsys` variable tells which operating system version of Epsilon is running, using the macros shown above and defined in `codes.h`.

```
short is_gui;
#define IS_NT      2      /* All Windows GUI versions */
```

The `is_gui` variable lets an EEL program determine if it’s running in a version of Epsilon that provides general-purpose dialogs. Currently, the variable is nonzero only in the Windows GUI version, where it is always set to `IS_NT`.

```
short is_unix;
#define IS_UNIX_TERM 1
#define IS_UNIX_XWIN 2
```

The `is_unix` variable is nonzero if Epsilon for Unix is running. It’s set to the constant `IS_UNIX_XWIN` if Epsilon is running as an X11 program, or `IS_UNIX_TERM` if Epsilon is running as a terminal program.

```
short is_unix_flavor;
#define IS_UNIX_LINUX 1
#define IS_UNIX_BSD   2
#define IS_UNIX_MACOS 3
```

The `is_unix_flavor` variable is nonzero if Epsilon for Unix is running. It’s set to the constant macro `IS_UNIX_LINUX` in the Linux version, `IS_UNIX_BSD` in the FreeBSD version, and `IS_UNIX_MACOS` in the Mac OS version.

```
short is_mswindows;
#define IS_MSWINDOWS_64BIT    1
#define IS_MSWINDOWS_32BIT    2
```

The `is_mswindows` variable is nonzero if a version of Epsilon for Windows is running, either the GUI version or the Windows console version. The constant `IS_MSWINDOWS_64BIT` means the 64-bit version of Epsilon is running. The constant `IS_MSWINDOWS_32BIT` means the 32-bit version of Epsilon is running.

```
int has_feature;
```

Epsilon provides the `has_feature` variable so an EEL function can determine which facilities are available in the current environment. Bits represent possible features. Often these indicate whether a certain primitive is implemented.

FEAT_ANYCOLOR	Epsilon can use all RGB colors, not just certain ones.
FEAT_GUI_DIALOGS	<code>display_dialog_box()</code> is implemented.
FEAT_FILE_DIALOG	<code>common_file_dlg()</code> is implemented.
FEAT_COLOR_DIALOG	<code>comm_dlg_color()</code> is implemented.
FEAT_SEARCH_DIALOG	<code>find_dialog()</code> is implemented.
FEAT_FONT_DIALOG	<code>windows_set_font()</code> is implemented.
FEAT_SET_WIN_CAPTION	<code>set_window_caption()</code> is implemented.
FEAT_OS_PRINTING	<code>print_window()</code> is implemented.
FEAT_WINHELP	<code>win_help_string()</code> and similar are implemented.
FEAT_WINHELP_NATIVE	The WinHelp program is always available.
FEAT_OS_MENUS	<code>win_load_menu()</code> and similar are implemented.
FEAT_ANSI_CHARS	Does this system normally use ANSI fonts, not DOS/OEM?
FEAT_EEL_RESIZE_SCREEN	Does EEL code control resizing the screen?
FEAT_INTERNET	Are Epsilon's Internet functions available?
FEAT_SET_FONT	Can EEL set the font via variables?
FEAT_MULT_CONCUR	Does Epsilon support multiple concurrent processes?
FEAT_DETECT_CONCUR_WAIT	Can Epsilon learn that a concurrent process waits for input?
FEAT_EEL_COMPILE	<code>eel_compile()</code> is implemented.
FEAT_LCS_PRIMITIVES	<code>lcs()</code> and related are implemented.
FEAT_PROC_SEND_TEXT	<code>process_send_text()</code> is implemented.
FEAT_UNICODE	16-bit Unicode characters are supported.

Figure 10.1: Bits in the `has-feature` variable.

The `FEAT_WINHELP` and `FEAT_WINHELP_NATIVE` bits differ on systems that don't include WinHelp support by default, but have it as an installable option (Windows Vista through Windows 8.1). For these, the first bit but not the second is present. Starting in Windows 10, neither bit is present.

```
#define MAX_CHAR    65535
```

The `MAX_CHAR` macro indicates the largest character code that can appear in a buffer.

```
ding()
maybe_ding(int want)    /* disp.e */
```

```

user int want_bell;           /* EEL variable */
user short beep_duration;
user short beep_frequency;

```

The `ding()` primitive produces a beeping sound, usually called the bell. It is useful for alerting the user to some error. Instead of calling `ding()` directly, however, EEL commands should call the `maybe_ding()` subroutine defined in `disp.e` instead. It calls `ding()` only if the variable `want_bell` is nonzero, and its parameter is nonzero. Pass one of the `bell_on_` variables listed on page 116 as the parameter. The sound that `ding()` makes is controlled by the `beep-duration` and `beep-frequency` variables. See page 116.

```

int clipboard_available()
int buffer_to_clipboard(int buffer_number, int flags,
                        int clipboard_format)
int clipboard_to_buffer(int buffer_number, int flags,
                        int clipboard_format)
#define CLIP_CONVERT_NEWLINES 1
#define CLIP_ADD_FORMAT      2
copy_line_to_clipboard(char *line, int flags)

```

The `clipboard_available()` primitive tells whether Epsilon can access the system clipboard in this environment. It returns nonzero if the clipboard is available, or zero if not. Epsilon for Windows can always access the clipboard. Epsilon for Unix can access the clipboard when it runs as an X11 program.

The `buffer_to_clipboard()` primitive copies the indicated buffer to the clipboard. A `clipboard_format` of zero means use the default format; otherwise, it specifies a particular Windows clipboard format code. If you provide `CLIP_CONVERT_NEWLINES` in the `flags` argument, Epsilon will add a `<Return>` character before each `<Newline>` character it puts on the clipboard. This is the normal format for clipboard text. Without this flag, Epsilon will put an exact copy of the buffer on the clipboard. With the `CLIP_ADD_FORMAT` flag, Epsilon will add the specified data to the clipboard without clearing its current contents first. The new data should use a different format code than the clipboard's current contents, and the additional format will be added. Only Epsilon for Windows recognizes this flag.

The `clipboard_to_buffer()` primitive replaces the contents of the given buffer with the text on the clipboard. The `clipboard_format` parameter has the same meaning as above. If `CLIP_CONVERT_NEWLINES` is used, Epsilon will strip all `<Return>` characters from the clipboard text before putting it in the buffer.

The `copy_line_to_clipboard()` subroutine copies a single line to the clipboard, also displaying it. Mode-specific functions called by the `copy-include-file-name` command often use it. Passing a flag of 1 makes it add a newline after the provided text; a flag of 2 makes it skip displaying a message indicating what it copied.

```

signal_suspend()

```

In Epsilon for Unix, the `signal_suspend()` primitive suspends Epsilon's job. Use the shell's `fg` command to resume it. When Epsilon runs as an X11 program, this primitive minimizes Epsilon instead.

10.4.2 Window System Primitives

```

windows_maximize()
windows_minimize()
windows_restore()

```

```

windows_foreground()
int windows_state()
int screen_to_window_id(int screen)

```

In Epsilon for Windows, and in Unix under X11, the `windows_maximize()`, `windows_minimize()`, and `windows_restore()` primitives perform the indicated action on the main Epsilon screen. The `windows_foreground()` primitive tries to make Epsilon the foreground window. (Under X11, some window managers may not let Epsilon do this. Also see the `server-raises-window` variable.)

The `windows_state()` primitive returns a code indicating the state of Epsilon's main window. The value `WINSTATE_MINIMIZED` indicates the window has been minimized or iconified. The value `WINSTATE_MAXIMIZED` indicates the window has been maximized. Zero indicates the window is in some other state.

The `screen_to_window_id()` primitive returns the system's window id (for X11) or window handle (for Windows) corresponding to a particular screen number. For Windows, specify a screen number of -1 to retrieve the window handle of the frame window that contains Epsilon's menu bar, title bar and so forth. It returns 0 if there is no screen with that number.

```

int drag_drop_result(char *file)
drag_drop_handler()
do_resume_client()
short reject_client_connections;

```

Epsilon uses the `drag_drop_result()` primitive to retrieve the names of files dropped on an Epsilon window using drag and drop, after receiving the event key `WIN_DRAG_DROP`. Pass the primitive a character array big enough to hold a file name. The primitive will return a nonzero value and fill the array with the first file name. Call the primitive again to retrieve the next file name. When the function returns zero, there are no more file names.

Epsilon uses this same method to retrieve server messages or DDE messages. When such a message arrives from another program, Epsilon parses the message as if it were a command line and then adds each file name to its list of drag-drop results. Epsilon for Unix doesn't support file drag and drop or DDE, only server messages from another copy of Epsilon.

When Epsilon returns the `WIN_DRAG_DROP` key, it also sets some mouse variables to indicate the source of the files that can be retrieved through `drag_drop_result()`. It sets `mouse_screen`, `mouse_x`, `mouse_y`, and similar variables to indicate exactly where the files were dropped. If the message arrived via DDE or due to `-add` or `-wait`, then `mouse_screen` will be -1.

The `drag_drop_result()` primitive returns 2 to indicate `-wait` was used to send the file name; 1 otherwise. If `-wait` was used in a client instance of Epsilon, the `do_resume_client()` primitive may be used to signal waiting clients that the user has finished editing the desired file and they may now resume.

The `drag_drop_handler()` subroutine in `mouse.e` handles the `WIN_DRAG_DROP` key. Don't bind this key to a subroutine with a different name; Epsilon requires that the `WIN_DRAG_DROP` key be bound to a function named `drag_drop_handler()` for correct handling of drag-drop.

A function may set the `reject_client_connections` variable to keep Epsilon from accepting any files or other messages from other clients, via either server messages or DDE. (Files may still be dropped on Epsilon.) The 1 bit keeps Epsilon from accepting these messages. Other instances of Epsilon that use the `-add` flag will not see an instance of Epsilon where this bit has been set.

The 2 bit in `reject_client_connections` lets Epsilon accept and queue such messages, but doesn't deliver them. Epsilon won't return the `WIN_DRAG_DROP` key as long as this bit is set, but will remember the list of queued files and deliver them once this bit has been cleared.

```
int dde_open(char *server, char *topic)
int dde_execute(int conv, char *msg, int timeout)
int dde_close(int conv)
```

Epsilon provides some primitives that you can use to send a DDE Execute message to another program under Windows.

First call `dde_open()` to open a conversation, providing the name of a DDE server and the topic name. It returns a conversation handle, or 0 if it couldn't open the conversation for any reason.

To send each DDE message, call `dde_execute()`. Pass the conversation handle from `dde_open()`, the DDE Execute message text to send, and a timeout value in milliseconds (10000, the recommended value, waits 10 seconds for a response). The primitive returns nonzero if it successfully sent the message.

Finally, call `dde_close()` when you've completed sending DDE Execute messages, passing the conversation handle. It returns nonzero if it successfully closed the connection.

WinHelp Interface

These Windows-only functions support WinHelp files, on those older Windows versions that can display them.

```
int win_help_contents(char *file)
```

The `win_help_contents()` primitive displays the contents page of the specified Windows help file. If the `file` parameter is "", it uses Epsilon's help file, displaying help on Epsilon. The function returns a nonzero value if it was successful.

```
int win_help_string(char *file, char *key)
```

The `win_help_string()` primitive looks up the entry for `key` in the specified Windows help file. If the `key` parameter is "", it shows the list of possible keywords. If the `file` parameter is "", it uses Epsilon's help file, displaying help on Epsilon. The function returns a nonzero value if it was successful.

```
windows_help_from(char *file, int show_contents)
```

The `windows_help_from()` subroutine wraps the above two subroutines. If there's a suitable highlighted region, it calls `win_help_string()` to display help on the keyword text in the highlighted region. Otherwise, it either displays the help file's contents topic (if `show_contents` is nonzero), or the help file's keyword index. The `windows_help_from()` subroutine also handles tasks like displaying an error if the user isn't running Epsilon for Windows.

The Menu Bar

```
int win_load_menu(char *file)
win_display_menu(int show)
```

The `win_load_menu()` primitive makes Epsilon read the specified menu file (normally `gui.mnu`), replacing all previous menu definitions. See the comments in the `gui.mnu` file for details on its format. The `win_display_menu()` primitive makes Epsilon display its menu bar, when its `show` parameter is nonzero. When `show` is zero, the primitive makes Epsilon remove the menu bar from the screen.


```
int win_menu_popup(char *menu_name)
```

The `win_menu_popup()` primitive pops up a context menu, as typically displayed by the right mouse button. The menu name must match one of the menu tags defined in the file `gui.mnu`, usually the tag `"_popup"`.

```
invoke_menu(int letter)
```

The `invoke_menu()` primitive acts like typing `Alt-letter` in a normal Windows program. For example, `invoke_menu('e')` pulls down the Edit menu. `Invoke_menu(' ')` pulls down the system menu. And `invoke_menu(0)` highlights the first menu item, but doesn't pull it down, like tapping and releasing the Alt key in a typical Windows program. (Also see the variable `alt-invokes-menu`.)

The Tool Bar

```
toolbar_create()
toolbar_destroy()
toolbar_add_separator()
toolbar_add_button(char *icon, char *help, char *cmd)
```

Several primitives let you manipulate the tool bar. They only operate in the Windows GUI version. The `toolbar_create()` primitive creates a new, empty tool bar. The `toolbar_destroy()` primitive hides the tool bar, deleting its contents. The `toolbar_add_separator()` primitive adds a blank space between buttons to the end of the tool bar.

The `toolbar_add_button()` primitive adds a new button to the end of the tool bar. The `cmd` parameter contains the name of an EEL function to run. The `help` parameter says what "tool tip" help text to display, if the user positions the mouse cursor over the button. The `icon` parameter specifies which icon to use. In this version, it must be one of these standard names:

STD_CUT	STD_PRINTPRE	VIEW_DETAILS
STD_COPY	STD_PROPERTIES	VIEW_SORTNAME
STD_PASTE	STD_HELP	VIEW_SORTSIZE
STD_UNDO	STD_FIND	VIEW_SORTDATE
STD_REDO	STD_REPLACE	VIEW_SORTTYPE
STD_DELETE	STD_PRINT	VIEW_PARENTFOLDER
STD_FILENEW	VIEW_LARGEICONS	VIEW_NETCONNECT
STD_FILEOPEN	VIEW_SMALLICONS	VIEW_NETDISCONNECT
STD_FILESAVE	VIEW_LIST	VIEW_NEWFOLDER
HIST_BACK	HIST_FORWARD	HIST_FAVORITES
HIST_VIEWTREE	HIST_ADDTOFAVORITES	

Run the commands `show-standard-bitmaps` or `show-view-bitmaps` to see what they look like. Run the command `standard-toolbar` to restore the original tool bar.

```
user char want_toolbar;
```

Epsilon uses the `want_toolbar` primitive variable to remember if the user wants a tool bar displayed, in versions of Epsilon which support this.

Printing Primitives

```

struct print_options {
    int flags;        // Flags: see below.
    int frompage;     // The range of pages to print.
    int topage;
    int height;
    int width;
};

/* Epsilon supports these printer flags. */
#define PD_SELECTION      0x00000001
#define PD_PAGENUMS      0x00000002
#define PD_PRINTSETUP     0x00000040

short select_printer(struct print_options *p)
page_setup_dialog()

```

In the Windows version of Epsilon, the `select_printer()` primitive displays a dialog box that lets the user choose a printer, select page numbers, and so forth. The flags and parameters are a subset of those of the Windows API function `PrintDlg()`. The primitive returns zero if the user canceled printing, or nonzero if the user now wants to print. In the latter case, Epsilon will have filled in the `height` and `width` parameters of the provided structure with the number of characters that can fit on a page of text using the selected printer.

The `page_setup_dialog()` displays the standard Windows page setup dialog, which you can use to set printer margins or switch to a different printer.

```

short start_print_job(char *jobname)
short print_eject()
short end_print_job()

```

After using the `select_printer()` primitive, an EEL program that wishes to print must execute the `start_print_job()` primitive. It takes a string specifying the name of this job in the print queue. The EEL program can then print one or more pages, ending each page with a call to `print_eject()`. After all pages have been printed, the EEL program must call `end_print_job()`.

```

short print_line(char *str, ?int scheme)
short print_window(int win)
int create_invisible_window(int width, int height, int buf)

```

To actually produce output, two primitives are available. The `print_line()` primitive simply prints the given line of text, and advances to the next line. It prints using the “text” color class in the current color scheme. If the optional parameter `scheme` is nonzero, Epsilon uses that color scheme instead.

The `print_window()` primitive prints the contents of a special kind of Epsilon window. The window must have been created by calling `create_invisible_window()`, passing it the desired dimensions of the window, in characters, and the buffer it should display. The `create_invisible_window()` primitive returns a window handle which can be passed to `print_window()`. An EEL program can move through the buffer, letting different parts of the buffer “show” in this window, to accomplish printing the entire buffer. The invisible window may be deleted using the `window_kill()` primitive once the desired text has been printed.

10.4.3 Timing

```
int time_ms()
time_begin(TIMER *t, int len)
int time_done(TIMER *t)
int time_remaining(TIMER *t)
```

The `time_ms()` primitive returns the time in milliseconds since some arbitrary event in the past. Eventually, the value resets to 0, but just when this occurs varies with the environment. In some cases, the returned value resets to 0 once a day, while others only wrap around after longer periods.

The `time_begin()` and `time_done()` primitives provide easier ways to time events. Both use the `TIMER` data type, which is built into Epsilon. The `time_begin()` primitive takes a pointer to a `TIMER` structure and a delay in hundredths of a second. It starts a timer contained in the `TIMER` structure. The `time_done()` primitive takes a pointer to a `TIMER` that has previously been passed to `time_begin()` and returns nonzero if and only if the indicated delay has elapsed. The `time_remaining()` primitive returns the number of hundredths of a second until the delay of the provided timer elapses. If the delay has already elapsed, the function returns zero. You can pass `-1` to `time_begin()` to create a timer that will never expire; `time_remaining()` will always return a large number for such a timer, and `time_done()` will always return zero.

Also see the `delay()` primitive on page 535.

```
current_time_and_date(char *s)
```

The `current_time_and_date()` subroutine fills in the string `s` with the current time and date. It uses the format specified by the `date-format` variable.

```
struct time_info {
    short year;
    short month;    /* 1-12 */
    short day;      /* 1-31 */
    short hour;     /* 0-23 */
    short minute;   /* 0-59 */
    short second;   /* 0-59 */
    short hundredth; /* 0-99 */
    short day_of_week; /* 0=Sunday ... 6=Saturday */
};
time_and_day(struct time_info *t_info)
```

The `time_and_day()` primitive requests the current time and day from the operating system, and fills in the `time_info` structure defined above. The structure declaration also appears in `eel.h`.

Notice that the `time_and_day()` primitive takes a *pointer* to a structure, not the structure itself. Here is an example command that prints out the time and date in the echo area.

```
#include "eel.h"

command what_time()
{
    struct time_info ts;
```

```

    time_and_day(&ts);
    say("It's %d:%d on %d/%d/%d.", ts.hour, ts.minute,
                                              ts.month, ts.day, ts.year);
}

```

10.4.4 Calling DLLs (Windows Only)

```

int call_dll(char *dll_name, char *func_name,
             char *ftype, char *args, ...)

```

The `call_dll()` primitive calls a function in a Windows DLL. The 32-bit version of Epsilon can only call 32-bit DLLs, and the 64-bit version can only call 64-bit DLLs. The `dll_name` parameter specifies the DLL file name. The `func_name` parameter specifies the name of the particular function you want to call.

The `ftype` parameter specifies the routine's calling convention. The character `C` specifies the C calling convention, while `P` specifies the Pascal calling convention. Most Windows DLLs use the Pascal calling convention, but any function that accepts a variable number of parameters must use the C calling convention.

The `args` parameter specifies the type of each remaining parameter. Each letter in `args` specifies the type of one parameter, according to the following table.

Character	Description
L	int INT, UINT, HWND, most other handles
I	int32 32-bit integer
S	far char * LPSTR
P	far void * LPVOID
R	far void ** LPVOID *

The codes L, S, P, and R represent 64-bit parameters (but 32-bit when calling 32-bit DLLs). The I character always represents a 32-bit parameter.

S represents a null-terminated string being sent to the DLL. P is passed similarly, but Epsilon will not check the string for null termination. It's useful when the string is an output parameter of the DLL, and may not be null-terminated before the call, or when passing structure pointers to a DLL.

R indicates that a DLL function returns a pointer by reference. Epsilon will pass the pointer you supply (if any) and retrieve the result. Use this for DLL functions that require a pointer to a pointer, and pass the address of any EEL variable whose type is "pointer to ..." (other than "pointer to function").

Here's an example, using `call_dll()` to determine the main Windows directory:

```

#define GetWindowsDirectory(dir, size) \
    (call_dll("kernel32.dll", "GetWindowsDirectoryW", \
              "p", "pi", dir, size))

char dir[FNAMELEN];

GetWindowsDirectory(dir, FNAMELEN);
say("The Windows directory is %s", dir);

```

After you call a function in a DLL, Epsilon keeps the DLL loaded to make future calls fast. You can unload a DLL loaded by `call_dll()` by including just the name of the DLL, and omitting the name of any function or parameters. For example, `call_dll("extras.dll");` unloads a DLL named `extras.dll`.

```
char *make_pointer(int value)
```

The `make_pointer()` primitive can be useful when interacting with system DLLs. It takes a machine address as a number, and returns an EEL pointer that may be used to access memory at that address. No error checking will be done on the validity of the pointer.

10.4.5 Running a Process

The following sections describe the EEL primitives for running a program. Here's a summary of the key differences among them.

	shell	concur_shell	pipe_text	winexec	run_viewer
Windows	•	•	•	•	•
Mac/Unix	•	•	•		•
Can wait	•		•	•	
Searches path	•	•	•		•
Opens docs	Win	Win	Win		•
Captures output	•	•	•		
Provides input		•	•		
Interactive		•			
Cmd line flags	•	•	•	•	
Hide/max/min	•			•	

Key:

Windows: Supported under Windows. **Mac/Unix:** Supported under macOS, Linux, and FreeBSD. **Can wait:** Primitive can be set to wait until subprocess has terminated before returning. **Searches path:** Primitive looks for executable along PATH environment variable, so executable's full path may be omitted.

Opens docs: You can provide the full path of a non-executable file instead of a program, and the appropriate program to open or display it will be run, per system settings. (Win indicates this only works under Windows.)

Captures output: Can retrieve text sent by a command-line program to its standard output into a buffer.

Provides input: Can send a block of text from a buffer to a command-line program's standard input.

Interactive: Can send text to a command-line program as it runs, interspersed with retrieving its output, allowing programmatic interaction.

Cmd line flags: Can run program and pass it parameters on its command line. **Hide/max/min:** Can set the GUI process's window hidden, or maximize or minimize it.

```
int shell(char *program, char *cline, char *buf, ?int flags)
```

The `shell()` primitive takes the name of an executable file (a program) and a command line, pushes to the program, and gives it that command line. The primitive returns the result code of the `wait()` system call, or `-1` if an error occurred. In the latter case, the error number is in `errno`.

The first argument to `shell()` is the name of the actual file a program is in, including any directory prefix. The second argument to `shell()` is the command line to pass to the program.

If the first argument to `shell()` is an empty string "", Epsilon behaves differently. In this case, Epsilon runs the appropriate shell command processor. (Note that "" is not the same as NULL, a pointer whose value is 0.) If the second argument is also "", Epsilon runs the shell interactively, so that it prompts for

commands. Otherwise, Epsilon makes the shell run only the command line specified in the second argument. Epsilon knows what flags to provide to the various standard shells to make them run interactively, or execute a single command and return, but you can set these if necessary. You can also set the command processor Epsilon should use. See page 150.

Under Windows, when you provide a nonempty first argument, Epsilon won't search the path for the specified file. To run a file on the path, put its name as the second argument and leave the first as "". This technique is also necessary to execute batch files, use internal commands like "dir", or do command-line redirection.

The third argument to `shell()` controls whether the output of the program is to be captured. If "", no capturing takes place. Otherwise the output is inserted in the specified buffer, replacing its previous contents.

In the Windows GUI version, and when Epsilon for Unix runs as an X11 program, Epsilon starts the program and then immediately continues without waiting for it to finish, whenever the first three arguments to `shell()` are "". Otherwise, Epsilon waits for the program to finish. The `SHELL_SYNCH` flag forces Epsilon to wait; the `SHELL_NO_SYNCH` flag tells Epsilon not to wait for the program.

The remaining flags for `shell()` only apply to the Windows version. `SHELL_HIDE` makes the resulting program's main window hidden; `SHELL_MINIMIZED` minimizes it, and `SHELL_MAXIMIZED` maximizes it. Normally Epsilon inserts an `EPSRUNS=Y` setting into the environment passed to the child process, in case some program wants to know if Epsilon invoked it. The `SHELL_KEEP_ENV` flag prevents that.

```
int do_push(char *cmdline, int cap, int show)
```

The `do_push()` subroutine is a convenient way to call `shell()`. It uses the command processor to execute a command line (so the command line may contain redirection characters and the like). If `cap` is nonzero, the subroutine will capture the output of the command to the process buffer. If `show` is nonzero, the subroutine will arrange to show the output to the user. How it does this depends on `cap`. To show captured output, Epsilon displays the process buffer after the program finishes. To show non-captured output, Epsilon (non-GUI versions only) waits for the user to press a key after the program finishes, before restoring Epsilon's screen. If `show` is -1, Epsilon skips this step.

This subroutine interprets the variable `start-process-in-buffer-directory` and takes care of displaying an error to the user if the process couldn't be run.

Concurrent Process Primitives

```
int concur_shell(char *program, char *cline,
                 ?char *curdir, char *buf, int flags)
short another_process();
int is_process_buffer(int buf)
```

The `concur_shell()` primitive also takes a program and a command line, with the same rules as the `shell()` primitive. It starts a concurrent process, with input and output connected to the buffer "process", just like the `start-process` command described on page 150 does. If you specify a buffer `buf`, it starts the process in that buffer. (Some versions of Epsilon support only one process buffer; in them the buffer name, if specified, must be "process".) If you specify a directory name in `curdir`, Epsilon starts the process with that current directory. The primitive returns 0 if it could start the process. If it couldn't, it returns an error code.

Normally Epsilon sets certain environment variables in the subprocess it creates, such as `EPSRUNS=C`. The `SHELL_KEEP_ENV` flag prevents that.

On Unix-based systems, Epsilon scans concurrent process output for a special command sequence set up by a shell script the `start-process` command normally invokes; this sends current directory info to

Epsilon. Internally the command sequence is used to set the `process-current-directory` variable and deleted before Epsilon inserts text in the buffer. The flag `SHELL_NO_CD_MAGIC` disables this processing.

Epsilon only receives concurrent process output and sends it input when Epsilon is waiting for you to press a key (or during a `delay()`—see page 535), but the process otherwise runs independently.

The `another_process()` primitive returns the number of active concurrent processes.

The `is_process_buffer()` primitive returns `ISPROC_CONCUR` if the specified buffer holds an active concurrent process, `ISPROC_PIPE` if the `buf_pipe_text()` primitive is sending output into it, or 0 if no concurrent process is associated with that buffer.

```
user buffer int type_point;
```

Characters from the process go into the process buffer at a certain position that we call the *type point*. The `type_point` variable stores this position.

When a process tries to read a character of input, Epsilon stops the process until there is at least one character following the type point, and when the process tries to read a line of input, Epsilon does not run the process until a newline appears in the section of the buffer after the type point. When a concurrent process is started by the `concur_shell()` primitive, the type point is initially set to the value of point in the specified buffer.

Internet commands for Telnet and FTP use `type_point` much like a process buffer does, to determine where to insert text into a buffer and where to read any text to be sent.

```
int process_input(?int buf)
#define PROCESS_INPUT_LINE 1
#define PROCESS_INPUT_CHAR 2
buffer int (*when_activity)();
concur_handler(int activity, int buf, int from, int to)
```

The `process_input()` primitive returns `PROCESS_INPUT_LINE` if the process is waiting for a character, `PROCESS_INPUT_CHAR` if the process is waiting for a line of input, and 0 if the process is running or there is no process. It operates on the buffer named “process” if no buffer number is specified.

Whenever Epsilon receives process output or sends it input, it calls an EEL function. The buffer-specific `when_activity` variable contains a function pointer to the function to call. If the variable is zero in a buffer, Epsilon won’t call any EEL function as it proceeds. For a typical process buffer, the `when_activity` variable points to the `concur_activity()` subroutine.

Just after a concurrent process inserts output in a process buffer, it calls this subroutine, passing `NET_RECV` as the activity. The `from` and `to` parameters mark the range of buffer text that was just received from the process. The `concur_activity()` subroutine responds to this message by coloring the inserted characters with the `color_class process_output` color, and similar tasks.

Epsilon calls this subroutine and passes `NET_SEND` when it detects that the concurrent process is now ready for input, and again as it sends the input to the process. When the process becomes ready for input, the subroutine will be called with a `from` parameter of zero. (Unix-based versions of Epsilon will only receive this notification when one of Epsilon’s shell scripts, such as `epsilon.sh`, has modified the prompt with a special code Epsilon interprets, and not when programs other than the shell are waiting for input.) When the process is sent a line of text, the subroutine will be called with a `from` of `PROCESS_INPUT_LINE`, and when the process is sent a single character it will be called with a `from` of `PROCESS_INPUT_CHAR`. In each case the `to` parameter will indicate the beginning of the input text (the value of `type_point` before the input begins).

Epsilon calls this subroutine and passes `NET_DONE` when the process exits. Its `from` parameter will hold the exit code, or 0 if Epsilon didn’t record this. Epsilon sets the buffer-specific `process_exit_status`

variable to the value `PROC_STATUS_RUNNING` when a process starts, and sets it to the process exit status (or 0) when the process exits.

Epsilon for Unix often cannot detect when a process is awaiting input. Therefore `process_input()` always returns zero, and a `NET_SEND` activity will typically not be signaled with a `from` parameter of zero.

```
int process_send_text(int buf, char *text, int len)
```

Normally input to a process running in a concurrent process buffer comes from text the user inserts into the buffer. The `process_send_text()` primitive provides a way to send text directly to the process, bypassing the buffer. This is especially useful for passwords, since if a password appears in the buffer it might be seen, or retrieved with `undo`. The primitive sends `len` characters from `text` to the process associated with the buffer `buf`.

The `FEAT_PROC_SEND_TEXT` bit of the `has_feature` variable indicates when this primitive is available.

```
int halt_process(?int hard_kill, int buf)
```

The `halt_process()` primitive has the same function as the `stop-process` command. A value of 0 for `hard_kill` makes the primitive act the same as `stop-process` with no argument. Otherwise, it is equivalent to `stop-process` with an argument. The function returns 1 if it succeeds, and 0 if it cannot signal the process for some reason. It operates on the buffer named “process” if no buffer number is specified.

```
int process_kill(?int buf)
```

The `process_kill()` primitive disconnects Epsilon from a running concurrent process, telling it to exit. The function returns 1 if it succeeds, and 0 if it cannot kill the process for some reason. It operates on the buffer named “process” if no buffer number is specified.

Other Process Primitives

```
int pipe_text(char *input, char *output, char *cmdline,
              char *curdir, int flags, int handler)
my_handler(int activity, int buf, int from, int to) // Sample.
int buf_pipe_text(int inputb, int outputb, char *cmdline,
                  char *curdir, int flags, ?int errorb)
```

The `pipe_text()` subroutine runs the program specified by `cmdline`, passing it the contents of a buffer as its standard input, and inserting its standard output into a second buffer (or the same buffer).

The input buffer name may be `NULL` if the process does not require any input. Epsilon provides a current directory of `curdir` to the process. It passes Epsilon’s current directory if `curdir` is `NULL` or `""`. This subroutine returns 0 and sets `errno` if the function could not be started, or returns 1 if the function started successfully.

The `PIPE_SYNCH` flag means don’t return from the subroutine until the process has finished. Without this flag, Epsilon starts the subprocess and then returns from `pipe_text()`, letting the subprocess run asynchronously.

The `PIPE_CLEAR_BUF` flag means empty the output buffer before inserting the process’s text (but do nothing if the process can’t be started); it’s convenient when the input and output buffers are the same, to filter a buffer in place.

The PIPE_NOREFRESH flag tells Epsilon not to refresh the screen each time more data is received from the process, and is most useful with PIPE_SYNC if you don't want the user to see the data until after it's been postprocessed in some way.

The PIPE_KEEP_ENV flag prevents Epsilon from modifying the environment it passes to the subprocess. By default, it passes settings such as EPSRUNS=P to the subprocess.

The PIPE_SKIP_SHELL flag makes Epsilon directly invoke the specified program, instead of using a shell as an intermediary. This results in improved performance, but command lines that use shell meta characters (like >file for redirection, | for pipelines, or file pattern wildcards) won't operate as desired. Only Epsilon for Unix supports this flag. When Epsilon prepares an argument list from the command line, it interprets and removes quotes which may surround arguments that contain spaces.

If handler is nonzero, it's the index of a function (that is, an EEL function pointer) to call each time text is received from the process, and when the process terminates. The handler function will be called with the buffer number into which more process output has just been inserted, and from and to set to indicate the new text. The parameter activity will be NET_RECV when characters have been received, or NET_DONE when the subprocess has exited. In the latter case from will hold the process exit code.

The command inserts text in a buffer at its type_point (see the previous section), the same as other functions that insert text into a buffer in the background.

Epsilon sets the buffer-specific process-exit-status variable in the output buffer to the value PROC_STATUS_RUNNING when a process starts, and sets it to the process exit status (or 0) when the process exits.

The pipe_text() subroutine described above is implemented using the buf_pipe_text() primitive. There are a few differences between these:

The buf_pipe_text() primitive uses buffer numbers, not buffer names. It won't create a buffer for you the way the subroutine will; the buffer must already exist. (Pass 0 for a buffer number if you don't need input.)

Instead of passing a function pointer for handler, you must instead set the buffer-specific when_activity variable in the output buffer prior to calling buf_pipe_text().

Pass a curdir of "", not NULL, to buf_pipe_text() to use Epsilon's current directory.

The buf_pipe_text() primitive accepts an additional, optional, parameter errorb. If nonzero, any output of the program sent to standard error will be sent to the errorb buffer instead of the outputb buffer. If errorb is zero, such output will appear in outputb along with standard output.

```
int winexec(char *prog, char *cmdline, int show, int flags)
/* Pass these values to winexec: */
#define SW_HIDE          0
#define SW_SHOWNORMAL    1
#define SW_SHOWMINIMIZED 2
#define SW_SHOWMAXIMIZED 3
#define SW_SHOWNOACTIVATE 4
#define SW_SHOW          5
#define SW_MINIMIZE      6
#define SW_SHOWMINNOACTIVE 7
#define SW_SHOWNA        8
#define SW_RESTORE       9
```

In Epsilon for Windows, the winexec() primitive runs a program, like the shell() primitive, but provides a different set of options. Normally, the second parameter to winexec() contains the command line to execute and the first parameter contains the full path of the program to execute.

The third parameter to `winexec()` specifies the window visibility state for the new program. It can be one of the values listed above.

The fourth parameter contains flag bits. The `SHELL_KEEP_ENV` flag prevents Epsilon from putting `EPSRUNS=Y` into the environment of the process it starts, as it does by default. The `SHELL_SYNCH` flag tells Epsilon to wait for the program to finish before returning from the `winexec()` primitive. By default, the primitive will return immediately.

This primitive returns the exit code of the program it ran. If an error prevented it from running the program, it returns `-1` and puts an error code in the global variable `errno`. When the primitive runs a program without waiting for it to finish, the primitive returns zero if the program started successfully.

```
int run_viewer(char *file, char *action, char *dir)
```

The `run_viewer()` primitive runs the program associated with the given file, using its Windows file association. The most common action is "Open", though a program may define others, such as "Print". The `dir` parameter specifies the current directory in which to run the program. The primitive returns nonzero if it was successful, or zero if it could not run the program or the program returned an error code. This primitive always returns zero in the Unix version of Epsilon, which uses a shell script `epsilon-viewer` to run a viewer.

10.5 Control Primitives

10.5.1 Control Flow

```
error(char *format, ...)
when_aborting()      /* control.e */
quick_abort()
```

Epsilon provides several primitives for altering the flow of control from one statement to the next. The `error()` primitive takes arguments like `say()`, displays the string as `say()` does, and then aborts the current command, returning to the main loop (see page 578). In addition this primitive discards any type-ahead and calls the user-defined subroutine `when_aborting()` if it exists. The standard version of `when_aborting()` optionally rings the bell and removes the erroneous command from any keyboard macro being defined. The primitive `quick_abort()` acts like `error()` but displays no message.

```
user char user_abort;
int abort_key;
check_abort()
```

The variable `user_abort` is normally 0. It is set to 1 when you press the key whose value is `abort_key`. To disable the abort key, set `abort_key` to `-1`. By default, the `abort_key` variable is set to Control-G. Use the `set-abort-key` command to set the `abort_key` variable. See page 106.

The primitive `check_abort()` calls `error()` with the argument "Canceled." if the variable `user_abort` is nonzero. Use the primitive `check_abort()` whenever a command can be safely aborted, since otherwise an abort will only happen when the command returns. Epsilon calls `check_abort()` internally during any searching operation (see page 446), when you use the `delay()` primitive (described below) to wait, or (optionally) during certain file matching primitives (see page 572) and file input/output (see page 497).

```
leave(?int exitcode)
when_exiting()      /* EEL subroutine */
```

The primitive `leave()` exits Epsilon with the specified exit code (or 0 if omitted). A nonzero exit code keeps Epsilon from saving any settings it normally would, such as the currently selected font's name, or the sizes of any resized dialogs.

Just before calling `leave()`, Epsilon's standard commands call any subroutine whose name starts with `do_when_exiting_`. It receives one integer parameter, nonzero if the user said to exit without checking for unsaved buffers or saving the session. It should return no result. (It can abort if it needs to prevent Epsilon from exiting; it should never do this if its parameter was nonzero, though.) Epsilon also calls the `when_exiting()` subroutine; modifying it was an earlier way to customize Epsilon's behavior when exiting.

```
delay(int hundredths, int condition, ?int buf)
```

The `delay()` primitive takes an argument specifying a period of time, in hundredths of a second, and a bit pattern specifying additional conditions (with codes specified in `codes.h`). It waits until one of the conditions occurs, or until the specified time limit is reached. A time limit of -1 means to wait forever.

The condition code `COND_KEY` makes Epsilon return when a key is pressed or any key-generating input event occurs (like a mouse event, or getting the focus). The condition code `COND_TRUE_KEY` is similar, but only returns on actual keys, not mouse events or other events. The condition code `COND_PROC` makes Epsilon return when a concurrent process is waiting for input, or has exited. The condition code `COND_PROC_EXIT` makes Epsilon return when a concurrent process has exited. For the last two conditions, Epsilon checks on the buffer specified by the optional parameter `buf`. If `buf` is missing or zero, it checks the buffer named "process". These conditions are ignored if no process is running in the specified buffer.

The condition flag `COND_RETURN_ABORT`, in combination with `COND_KEY`, makes the `delay()` primitive return if the user presses the abort key, instead of aborting by calling the `check_abort()` primitive. (Note that if you don't specify `COND_KEY` or `COND_TRUE_KEY` as well, the primitive ignores all keys, including the abort key.)

This function varies a bit from one operating system to another. For example, the Unix version of Epsilon can't detect when a process is currently waiting for input, so it can only return when a process exits. Also see the timing functions on page 527.

```
int do_recursion()
leave_recursion(int val)
int recursive_edit() /* control.e */
int recursive_edit_preserve() /* control.e */
char _recursion_level;
```

The `do_recursion()` primitive starts a new loop for getting characters and interpreting them as commands. A recursive edit preserves the current values of the variables `has_arg`, `iter`, `this_cmd`, and `prev_cmd`, but does not preserve the current buffer, window, or anything else. (See page 578.) Exit the recursion by calling the `leave_recursion()` primitive. It arranges for the main loop to exit, instead of waiting for another key to be executed. The call to `do_recursion()` will then return with a value of `val`, the argument of the call to `leave_recursion()`.

Sometimes a recursive edit is done "secretly," and the user doesn't know that one is being used. For example, when Epsilon reads the name of a file using completion, it's actually doing a recursive edit. Keys like `<Space>` exit the recursive edit with a special code, and the function that did the recursive edit displays a menu, or whatever is needed, and then does another recursive edit.

Other times (typing `Ctrl-r` in `query-replace`, for example), the user is supposed to exit the recursive edit explicitly using the `exit-level` command. When you're supposed to use `exit-level` to exit, Epsilon displays

extra []'s in the mode line as a reminder. The `recursive_edit()` subroutine does a recursive edit, and arranges for these []'s to appear by modifying the `_recursion_level` variable. It contains the number of extra []'s to display. The `recursive_edit()` subroutine returns the value returned by `do_recursion()`.

The `recursive_edit_preserve()` subroutine calls `recursive_edit()`. If the user changes the current buffer or window during the recursion, `recursive_edit_preserve()` returns to the original buffer or window before itself returning. If the user deleted the original buffer or window during the recursive edit, this subroutine remains in the new buffer or window and returns 0. It returns 1 otherwise.

If you call `leave_recursion()` when there has been no matching `do_recursion()`, Epsilon automatically invokes the command `exit`. If `exit` returns instead of calling the primitive `leave()`, Epsilon begins its main loop again.

```
int setjmp(jmp_buf *location)
longjmp(jmp_buf *location, int value)
```

Epsilon implements aborting by two special primitives that allow jumping from a function to another point in that function or any of the functions that called it. The `setjmp()` primitive marks the place to return, storing the location in a variable declared like this:

```
jmp_buf location;
```

After calling `setjmp()` with a pointer to this structure, you can return to this place in the code at any time until this function returns by calling the `longjmp()` primitive. The first argument is a pointer to the same structure, and the second argument may be any nonzero value.

The first time `setjmp()` is called, it returns a zero value. Each time `longjmp()` is called, Epsilon acts as if it is returning from the original `setjmp()` call again, returning the second argument from the `longjmp()`. For example:

```
one()
{
  jmp_buf location;

  if (setjmp(&location)){
    stuff("Back in one\n");
    return;
  } else
    stuff("Ready to go\n");
  two(&location);
}

two(loc)
jmp_buf *loc;
{
  stuff("In two\n");
  longjmp(loc, 1);
  stuff("Never get here\n");
}
```

This example inserts the lines

```
Ready to go
In two
Back in one
```

```
jmp_buf *top_level;
```

The `error()` primitive uses the jump buffer pointed to by the `top_level` variable. If you wish to get control when the user presses the abort key, temporarily change the value of `top_level` to refer to another jump buffer. Make sure you restore it, however, or subsequent aborting may not work.

10.5.2 Character Types

```
int isspace(int ch)
int isdigit(int ch)
int isalpha(int ch)
int islower(int ch)
int isupper(int ch)
int iscntrl(int ch)
int isgraph(int ch)
int ispunct(int ch)
int isprint(int ch)
int isxdigit(int ch)
int isalnum(int ch) /* basic.e */
int isident(int ch) /* basic.e */
int any_uppercase(char *p)
```

Epsilon has several primitives that are helpful for determining if a character is in a certain class. The `isspace()` primitive tells if its character argument is a space, tab, or newline character. It returns 1 if it is, otherwise 0.

In the same way, the `isdigit()` primitive tells if a character is a digit (one of the characters 0 through 9), and the `isalpha()` primitive tells if the character is a letter. The `islower()` and `isupper()` primitives tell if the character is a lower case letter or upper case letter, respectively.

The `iscntrl()` primitive tells if a character is a control character, `isgraph()` if a character is a graphical character (has a printed representation, not a space or a control character), `ispunct()` if a character is a punctuation character, `isprint()` if a character is printable (not a control character), and `isxdigit()` if a character is a hex digit.

The `isalnum()` subroutine returns nonzero if the specified character is alphanumeric: either a letter or a digit. The `isident()` subroutine returns nonzero if the specified character is an identifier character: a letter, a digit, or the `_` character.

All functions in this section also handle Unicode characters appropriately: the `isspace()` primitive, for instance, also returns 1 for characters with a Unicode category of Z, meaning separators, and primitives that report on lowercase letters understand any Unicode character that's a lowercase letter.

The `any_uppercase()` subroutine returns nonzero if there are any upper case characters in its string argument `p`.

```
int tolower(int ch)
int toupper(int ch)
```

The `tolower()` primitive converts an upper case letter to the corresponding lower case letter. It returns a character that is not an upper case letter unchanged. The `toupper()` primitive converts a lower case letter to its upper case equivalent, and leaves other characters unchanged.

```
int set_character_property(int ch, int propcode, int value)
```

You can alter the rules Epsilon uses for determining if a particular character is alphabetic, uppercase, or lowercase, and how Epsilon case-folds when searching, sorting or otherwise comparing text, using the `set_character_property()` primitive. It takes the numeric code of the character whose properties you want to modify, a property code indicating which of its properties to access, and a new value for that property.

The property code `CPROP_CTYPE` sets whether the `isalpha()`, `isupper()`, `islower()`, and `isdigit()` primitives consider a character alphabetic, uppercase, lowercase, or a digit, respectively. These attributes are independent, though there are conventions for their use. (For instance, only alpha characters generally have a case, no character is both uppercase and lowercase, and so forth.) The bits `C_ALPHA`, `C_LOWER`, `C_UPPER`, and `C_DIGIT` represent these attributes. The bits also control whether the regular expressions `<digit>`, `<alpha>`, `<alphanum>`, and `<word>` match these characters; see page 69.

Similarly, the bits `C_CNTRL`, `C_GRAPH`, `C_PUNCT`, and `C_XDIGIT` may be used to modify the `isctrl()`, `isgraph()`, `ispunct()`, and `isxdigit()` primitives, respectively, using `CPROP_CTYPE`.

The property code `CPROP_TOLOWER` controls what value the `tolower()` primitive returns for the specified character, and the property code `CPROP_TOUPPER` controls what value the `toupper()` primitive returns for it.

The property code `CPROP_FOLD` controls how Epsilon case-folds that character during searching, sorting, and similar functions, whenever case folding is in use. It specifies a replacement character to be used in place of the original during comparisons. The complete set of case-folding properties must follow two rules: if some character X folds to Y, then Y must fold to itself, and character codes below 256 must never fold to a value greater than or equal to 256. (If a particular group of characters should be treated as equal when searching, setting the case folding property of each to the code of the lowest-numbered one is sufficient to comply with these rules.)

The primitive returns the previous value of the specified property of that character. If the new value is out of range for the property (such as a negative value), it will be ignored, and the primitive will just return the current value. You can use this to retrieve the current properties of a character without changing them.

The special property code `CPROP_DISP_WIDTH` may be used to retrieve the width in columns of any character in the current font in Epsilon for Windows. It returns 1 for normal characters, 0 for zero-width characters, and 2 for those characters treated as doublewidth (if Epsilon's experimental doublewidth feature to better support Chinese, Japanese and Korean text has been enabled; it's disabled by default). It returns -1 if called during startup before Epsilon has selected a font, and in non-Windows versions (where all characters have a font display width of 1).

The special property code `CPROP_DEF_REPL` may be used to retrieve replacement Unicode characters for those in the range 128–255 in Epsilon for Windows using the current language settings. Some Extended ASCII characters in this range have different character numbers in Unicode, so Epsilon uses this function to map undisplayable characters in that range to displayable characters (for display purposes only). The function returns the original character for any values outside its range.

Epsilon doesn't store current character properties in its state file. If you want to use non-default properties all the time, write a startup function that calls this primitive. See page 552.

Epsilon always starts with character classifications based on standard Unicode properties, except for the Windows Console version. That version, when running with a DOS/OEM character set (see the `console-ansi-font` variable), begins with its classifications for 8-bit characters set to match the current OEM font.

```
int get_direction()          /* window.e */
```

The `get_direction()` subroutine converts the last key pressed into a direction. It understands arrow keys, as well as the equivalent control characters. It returns `BTOP`, `BBOTTOM`, `BLEFT`, `BRIGHT`, or -1 if the key doesn't correspond to any direction.

10.5.3 Examining Strings

```
int strlen(char *s)
```

Epsilon provides various functions for manipulating strings, or equivalently, zero-terminated arrays of characters. (General-purpose functions for modifying strings are covered in the next section.) The `strlen()` primitive returns the length of a string. That is, it tells the position in the array of the first zero character.

```
int strcmp(char *first, char *second)
int strncmp(char *first, char *second, int count)
```

The `strcmp()` primitive tells if two strings are identical. It returns 0 if all characters in them are the same (and if they have the same length). Otherwise, it returns a negative number if the lexicographic ordering of these strings would put the first before the second. It returns a positive number otherwise. The `strncmp()` primitive is like `strcmp()`, except only the first count characters matter.

```
int strcasecmp(char *first, char *second)
int strncasecmp(char *first, char *second, int count)
int charcasecmp(int first, int second)
```

Epsilon also has similar comparison primitives that consider upper case and lower case letters to be equal. The `strcasecmp()` primitive acts like `strcmp()` and the `strncasecmp()` primitive acts like `strncmp()`, but if the buffer-specific variable `case_fold` is nonzero, Epsilon folds characters in the same way searching or sorting would before making the comparison. The `charcasecmp()` primitive takes two characters and performs the same comparison on them. For characters *a* and *b*, `charcasecmp('a', 'b')` equals `strcmp("a", "b")`. (EEL also recognizes the corresponding ANSI C name `stricmp()` instead of `strcasecmp()`.)

```
int compare_chars(char *str1, char *str2, int num, int fold)
```

The `compare_chars()` primitive works like `strcmp()`, except that it makes no assumptions about zero-termination. It takes two strings and a size, then compares that many characters from each string. If the strings exactly match, `compare_chars()` returns zero. If `str1` would be alphabetically before `str2`, it returns a negative value. If `str2` would be alphabetically before `str1`, it returns a positive value. It ignores the case of the characters when comparing if `fold` is nonzero.

```
char *index(char *s, int ch)
char *rindex(char *s, int ch)
char *strstr(char *s, char *t)
char *strpbrk(char *s, char *charset)
char *strpbrk_cnt(char *s, char *charset, int skip)
```

The `index()` primitive tells if a character `ch` appears in the string `s`. It returns a pointer to the first appearance of `ch`, or a null pointer if there is none. The `rindex()` primitive works the same, but returns a pointer to the last appearance of `ch`. (EEL also recognizes the corresponding ANSI C names `strchr()` instead of `index()` and `strrchr()` instead of `rindex()`.)

The `strstr()` primitive searches the string `s` for a copy of the string `t`. It returns a pointer to the first appearance of `t`, or a null pointer if there is none. It case-folds as described above for `strcasecmp()`.

The `strpbrk()` subroutine returns a pointer to the first character in `s` that appears in the list of characters `charset`. Both strings must be null-terminated. If the strings have no characters in common, it returns a null pointer.

The `strpbrk_cnt()` subroutine is similar, but it skips over the first `skip` characters in `s` that also appear in `charset`. For instance, with `skip` set to 1, it returns a pointer to the second character in `s` that also appears in `charset`.

```
int fpatmatch(char *s, char *pat, int prefix, int flags)
#define FPAT_FOLD 1
#define FPAT_IGNORE_SQUARE_BRACKETS 2
```

The `fpatmatch()` primitive returns nonzero if a string `s` matches a pattern `pat`. It uses a simple filename-style pattern syntax: `*` matches any number of characters; `?` matches a single character, and `[a-z]` match a character class (with the same character class syntax as other patterns in Epsilon). It also recognizes `|` to permit alternatives. If `prefix` is nonzero, `s` must begin with text matching `pat`; otherwise `pat` must match all of `s`.

The `flags` parameter recognizes two bits. The `FPAT_FOLD` bit makes Epsilon fold characters before comparing, according to the current buffer's folding rules. The `FPAT_IGNORE_SQUARE_BRACKETS` bit makes Epsilon treat the character `[` in a pattern like any other, instead of interpreting it as the start of a character class.

```
int string_matches_regex(char *str, char *pat, int fold)
int string_matches_pattern(char *str, char *pat)
int regex_replace_in_string(char *dest, char *src, char *pat, char *repl)
```

The `string_matches_regex()` subroutine returns nonzero if the start of the given string matches the regular expression pattern. Use `<eof>` at the end of the pattern to check if the entire string matches. It does case-folding if `fold` is nonzero.

The similar `string_matches_pattern()` subroutine returns the length of match (which differs from the above only with patterns that can match zero-length text), and uses `case_fold.default`.

Both return zero when given an invalid regular expression pattern.

The `regex_replace_in_string()` subroutine copies `src` to `dest`, performing a regular expression replacement as it does. As in `string_replace()`, which it runs, the replacement text may contain `#` sequences to interpolate text from each match. It returns the number of replacements it performed, or -1 if an error like an invalid search pattern was detected. If `src` and `dest` are the same, the text is updated in place.

```
int word_in_list(char *word, char *list, int fold)
int starts_with_in_list(char *word, char *list, int fold)
```

The `word_in_list()` subroutine returns nonzero whenever the text in `word` appears in the `|`-separated list of words `list`. “Word” here means any text that doesn't contain an actual `|` character. The list of words must begin and end with `|` delimiters. The similar `starts_with_in_list()` subroutine returns nonzero whenever `word` starts with one of the words in the list. Both do case-folding if `fold` is nonzero. They are faster than the regular-expression-based subroutines above.

10.5.4 Modifying Strings

```
strcpy(char *tostr, char *fromstr)
strncpy(char *tostr, char *fromstr, int count)
copy_expanding(char *src, char **dest, int minlen)
```

The `strcpy()` primitive copies the null-terminated string `fromstr` to the array at `tostr`, including the terminating null character. The `strncpy()` primitive does the same, but always stops when `count` characters have been transferred, adding an additional null character to the string at `tostr` if necessary.

The `copy_expanding()` subroutine helps work with text that has no fixed length, stored in a dynamically allocated character pointer, not a fixed-length character array. Pass a pointer to a `char *` variable as `dest`, and the subroutine will resize it as needed to hold `src`. The `char *` variable may hold NULL initially. The `minlen` parameter provides a minimum allocation length for the result.

```
strcat(char *tostr, char *fromstr)
strncat(char *tostr, char *fromstr, int count)
```

The `strcat()` primitive concatenates (or appends) the string at `fromstr` after the string at `tostr`. For example, if `fromstr` points at the constant string “def” and `tostr` is an array of 10 characters that contains “abc” (and then, of course, a null character, plus 6 more characters with any value), then `strcat(tostr, fromstr)`; makes the array `tostr` contain “abcdef” followed by a null character and 3 unused characters.

The `strncat()` primitive works similarly. It appends at most `count` characters from `fromstr`, and ensures that the result is zero-terminated by adding a null character if necessary. Note that the `count` limits the number of characters appended, not the total number of characters in the string.

```
set_chars(char *ptr, char value, int count)
```

The `set_chars()` primitive sets all the `count` characters in a character array `ptr` to the given value.

```
int sprintf(char *dest, char *format, ...)
```

The `sprintf()` primitive is the most powerful string building primitive Epsilon provides. It takes two or more arguments. The first is a character array. The remaining arguments are in the format that `say()` uses: a format string possibly followed by more arguments. (See page 480.) Instead of printing the string that is built on the screen, it copies the string into the destination array, and returns the number of characters copied.

```
int expand_string_template(char *dest, char *template,
                           char *keys, char *vals[])
```

The `expand_string_template()` subroutine uses a template to construct a new string `dest`. The template contains zero or more escape sequences, each a percent character `%` followed by another letter. The subroutine replaces each escape sequence with its corresponding value. The allowed escape sequence characters should be listed in `keys`, and their replacement values should be listed in the array of strings `vals` in the same order. Epsilon will interpolate the sequences to construct `dest`, returning nonzero to indicate an error in the template (usually an unknown escape sequence).

For instance, if `keys` contains “ad”, and the supplied `vals` array has been set so `vals[0]` is “happy” and `vals[1]` is “tiger”, then the template “the %a is %d today” is copied to `dest` as “the tiger is happy today”, and the subroutine will return 0.

The subroutine provides two built-in codes. For `%x`, it substitutes the directory name that contains Epsilon’s executable. For `%X`, it substitutes the same, but (under Windows) converted to its 8.3 filename alias using the `convert_to_8_3_filename()` primitive. These built-in codes shouldn’t appear in the `keys` string.

10.5.5 Byte Arrays

These functions operate on 8-bit byte arrays, not 16-bit characters.

```
int memcmp(byte *str1, byte *str2, int num)
int memfcmp(byte *str1, byte *str2, int num)

memcpy(byte *tostr, byte *fromstr, int num)
memset(byte *ptr, char value, int count)
```

The `memcmp()` and `memfcmp()` primitives compare 8-bit bytes, not 16-bit characters like the `compare_chars()` and `strcmp()` primitives do. The `memfcmp()` primitive ignores case, `memcmp()` doesn't. They return values like the others, and don't stop comparing when an array element has a zero value, as `strcmp()` does.

The `memcpy()` primitive copies exactly `num` bytes from the second byte array to the first.

The `memset()` primitive sets all the `count` bytes in a byte array `ptr` to the given value.

```
chars_to_bytes(byte *b, char *s)
bytes_to_chars(char *s, byte *b)
```

The `chars_to_bytes()` function copies the null-terminated character string `s` to the byte array `b`, discarding the upper 8 bits of each 16-bit character.

The `bytes_to_chars()` function copies the null-terminated byte string `b` to the character array `s`.

10.5.6 Memory Allocation

```
char *malloc(int size)
char *realloc(char *ptr, int size)
free(char *ptr)
```

Epsilon maintains a pool of memory and provides primitives for allocating and deallocating blocks of any size. The `malloc()` primitive takes an `int` giving the number of characters of space required, and returns a pointer to a block of that size.

The `realloc()` primitive takes a pointer previously allocated with `malloc()`. First, it tries to expand the block to the requested size. If it cannot do that, it allocates another block of the requested size, then copies the old characters to the new block. In either case, it returns a pointer to a block of the requested size.

The `free()` primitive takes a pointer that `malloc()` previously returned and puts it back into the storage pool. Never use a block after you free it.

```
char *strsave(char *s)
char *strkeep(char *s)
```

For convenience, Epsilon provides a primitive to copy a string to an allocated block of the proper size. The `strsave()` primitive is used when a string needed later is stored in an array that must be reused. The primitive returns a pointer to the copy of the string it makes. The `free()` primitive may be given this pointer when the string is no longer needed.

The `strkeep()` subroutine also saves a string so it may be used later, returning a pointer to the copy. It's often used to store a mode name for use with the `major_mode` variable. Unlike `strsave()`, calling `strkeep()` on the same text multiple times always reuses the same saved block. Strings allocated by `strkeep()` may not be freed; they remain until Epsilon exits.

```
user int mem_in_use;
```

The `mem_in_use` variable gives the space in bytes Epsilon is now using for miscellaneous storage (not including buffer text).

```
set_swapname(char *path)
```

If Epsilon can't fit all your files in available memory, it will swap parts to disk. The parts are contained in one or more swap files. The `set_swapname()` primitive tells Epsilon what directories to use for swap files, if it needs them. The argument is a string containing a list of *directories* in which to place swap files, as described under the `-fs` command line flag. After swapping has begun, this primitive has no effect. Supplying an empty argument `"` makes Epsilon use the standard place for swapping, as described under the `-fs` command line switch on page 14.

10.5.7 The Name Table

```
int final_index()
```

Epsilon keeps track of all EEL variables, commands, subroutines, key tables, color schemes, and keyboard macros in its *name table*. Each of these items has an entry there that lists its name, type, value, and additional information. An EEL program can access the table using a numeric index, like an array index. The first valid index to the name table is 1, and the `final_index()` primitive returns the last valid index. The index is based on the order in which the names were defined.

All variables appear in the name table, including primitive variables. Primitive functions (like most of those defined in this chapter) and EEL's `#define` textual macros are not in the name table. A state file contains an exact copy of a name table (plus some additional information).

Each entry contains the name of the item, a type code, a debugging flag, a help file offset, and whatever information Epsilon needs internally to make use of the item. When executing an EEL program, Epsilon automatically uses the table to find the value of a variable, for example, or execute a command. You can manipulate the table with EEL functions.

```
int find_index(char *name)
```

There are two ways to get an index if you have the name of an item. The `find_index()` primitive takes an item name as a string and returns the index of that item, or 0 if there is no such item. If the item is an EEL command or subroutine, casting its function pointer to a short also yields the index. For example, `(short) forward_word` gives the index of the command `forward-word` if `forward_word()` has been declared previously in the source file the expression appears in.

```
char *name_name(int index)
int name_type(int index)      /* codes: */
#define NT_COMMAND      1    /* normal bytecode function */
#define NT_SUBR         2    /* hidden bytecode function */
#define NT_MACRO        3    /* keyboard macro */
#define NT_TABLE        4    /* key table */
#define NT_VAR          5    /* normal variable */
#define NT_BUFVAR       6    /* buffer-specific variable */
#define NT_WINVAR       7    /* window-specific variable */
#define NT_COLSCHEME    8    /* color scheme */
```

```
#define NT_BUILTVAR    9      /* built-in variable */
#define NT_AUTOLOAD   10     /* load cmd from file */
#define NT_AUTOSUBR   11     /* load subr from file */
```

The primitives `name_name()` and `name_type()` return the name and type of a table entry, respectively. They each take an index into the name table and return the desired information. The value returned by `name_name()` is only valid until the next call to this function. Copy the name if you want to preserve it.

The codes for `name_type()` are in the standard include file `codes.h`.

```
int try_calling(char *name)
```

The `try_calling()` primitive calls a subroutine or command if it exists and doesn't complain if the function does not exist. It takes the name of the function to call. It returns 0 if the function doesn't exist. The function it calls must not require arguments.

```
int call_with_arg_list(int func, char *argtypes, int intargs[],
                      char *strargs[], ?char **res)
```

The `call_with_arg_list()` primitive can call an EEL function that takes only integer and string parameters, specified by its name table index `func`. The caller must specify the types of the parameters to pass in `argtypes`, and supply the integer and string parameters in the arrays `intargs` and `strargs`, respectively. The `argtypes` value is a list of characters: "i" for an integer, or "s" for a string. For each "i", Epsilon uses the next entry in the `intargs` array, and for each "s", the `strargs` array.

If the EEL function returns an integer, the `call_with_arg_list()` primitive returns it. If it returns a character pointer, pass the address of a character pointer, and the primitive will fill it in with the return value.

The `func` parameter may refer to a keyboard macro if `argtypes` is empty.

```
int drop_name(char *name)
```

To delete an item from the name table, use the `drop_name()` primitive. It returns 0 if it deleted the name, 1 if there was no such name in the name table, and 2 if there was such a name but it couldn't be deleted because it is currently in use.

```
int replace_name(char *old, char *new)
```

The `replace_name()` primitive renames an item in the name table. It returns 0 if the name change was successful, 1 if the original name did not exist, and 2 if the name change was unsuccessful because another item had the new name already. Any references to the original item result in an error, unless you provide a new definition for it later.

Sometimes when writing an Epsilon extension, you may wish to redefine one of Epsilon's built-in subroutines (`getkey()`, for example) to do something in addition to its usual action. You can, of course, simply modify the definition of the function, adding whatever you want. Unfortunately, if someone else gives you an extension that modifies the same function, it will overwrite your version. You'll have the same problem when you get a new version of Epsilon—you'll have to merge your change by hand.

```
#define REPLACE_FUNC(ext, func) ....
/* definition omitted */
```

Alternatively, you can create an extension that modifies the existing version of a function, even if it's already been modified. The trick is to replace it with a function that calls the original function. This can be done from a `when_loading()` function by using the `replace_name()` and `drop_name()` primitives, but `eel.h` defines a macro that does all of this. The `REPLACE_FUNC()` macro takes the name of the extension you're writing, and the name of the existing subroutine you want to replace. It doesn't really matter what the extension name is, just so long as no other extension uses it.

Here's an example. Suppose you're writing an extension that displays "Hello, world" whenever you start Epsilon. You've decided to name the extension "hello", and you want Epsilon's `start_up()` function to do the work. Here's what you do:

```
new_hello_start_up()    /* will be renamed to start_up */
{
    say("Hello, world");
    hello_start_up(); /* call old (which will have this name) */
}

REPLACE_FUNC("hello", "start-up")
```

Notice the steps: first you have to define a function with a name of the form `new_<extension-name>_<replaced-function-name>`. Make sure it calls a function named `<extension-name>_<replaced-function-name>`. Then do the `REPLACE_FUNC()`, providing the two names. This will rename the current `<replaced-function-name>` to `<extension-name>_<replaced-function-name>`, then rename your function to `<replaced-function-name>`.

10.5.8 Built-in and User Variables

Variables that are automatically defined by Epsilon, and have no definition in `eel.h`, are called built-in variables. These include `point`, `bufnum`, and most of the primitive variables described in this chapter. All such built-in variables have entries in Epsilon's name table, so that you can see and set them using commands like `set-variable` or `set-any-variable`. Built-in variables have a `name_type()` code of `NT_BUILTVAR`.

```
int get_num_var(int i)
set_num_var(int i, int value)

char *get_str_var(int i)
set_str_var(int i, char *value)
```

Epsilon has several primitives that let you get and set the value of numeric and string global variables (including both built-in and ordinary, user-defined variables). Each primitive takes a name table index `i`. The `get_num_var()` and `get_str_var()` primitives return the numeric or string value (respectively) of the indicated variable, while the `set_num_var()` and `set_str_var()` primitives set the variable. If you provide an index that doesn't refer to a variable of the correct type, the setting functions do nothing, while the getting functions return zero. (See the `vartype()` primitive below.) The `set_str_var()` primitive only operates on variables with a character pointer data type, not on character arrays. Use `varptr()` below to modify character arrays.

The `set-variable` command and similar functions look for and try to call a function named `when_setting_varname()` after setting a variable named `varname`. For most variables a function with that name doesn't exist, and nothing happens. The `want-code-coloring` variable is an example of a variable

with a `when_setting()` function. Its `when_setting()` function sets various other variables to match `want-code-coloring`'s new value.

Any user attempts to set a variable (such as running `set-variable` or loading a command file) will call such a function, but an ordinary assignment statement in an EEL function will not. If you write an EEL function that sets a variable with a `when_setting()` function, you should call the function explicitly after setting the variable.

```
int name_user(int i)
set_name_user(int i, int is_user)
```

For each global variable, built-in or not, Epsilon records whether or not it is a “user” variable. Some commands such as `set-variable` only show user variables. Otherwise, Epsilon treats user variables the same as others. The `name_user()` primitive returns non-zero if the variable with the given name table index is a user variable, and the `set_name_user()` primitive sets whether a variable with a particular name table index is a user variable.

```
user int my_var;          // sample declaration
```

By default, variables you declare with EEL are all non-user variables, hidden from the user. If the user is supposed to set a variable directly in order to alter a command's behavior, put the `user` keyword before its global variable definition to make it a user variable. (In previous versions, Epsilon used a convention that any non-user variables you defined had to start with an underscore character, and all others were effectively user variables. This convention still works: `set-variable` will still exclude such variables from normal completion lists.)

```
int ptrlen(char *p, ?int in_bytes)
typedef struct eel_pointer {    /* format of EEL pointer */
    int base, size, value;
} EEL_PTR;
```

The `ptrlen()` primitive takes a pointer of any type and returns the size in characters of the object it points to. The value of `ptrlen(p)` is the lowest value `i` for which `((char *)p)[i]` is an illegal dereference. If its optional second argument is nonzero, it returns its count in bytes, not characters. (Characters are 16 bits wide, while bytes are 8 bits wide.)

The `EEL_PTR` type, defined in `lowlevel.h`, is a structure representing the internal format of an EEL pointer (except for function pointers, which are represented as short integers internally). An EEL pointer consists of a base, a size, and a value. The base and value are standard system pointers, and the size is an integer. Epsilon compares the three fields to catch invalid pointer usage.

Whenever a function dereferences a pointer, Epsilon checks that the fields are consistent. That is, it makes sure that `value` is greater than or equal to `base`, and that `value` is less than `base+size`. Epsilon will report an illegal dereference if these conditions are not met.

When Epsilon constructs a pointer, it sets the base field to the start of the block of storage within which the pointer points, and sets the size field to the size of the block of storage, in bytes. Epsilon then sets the value field to the actual address to which the pointer points. For example, if an EEL pointer `p` points to the letter ‘c’ in the string “abcd” (which is terminated by a null character), the size field of `p` will contain 10 (since five 16-bit characters require ten bytes), the base field will point to the ‘a’, and the value field will point to the ‘c’. Adding an integer to `p` will change only the value field. Notice that the modified version of `p` is “consistent” according to the rules above exactly when dereferencing it would be legal: `*(p - 2)`, `*(p - 1)`, `*p`, `*(p + 1)` and `*(p + 2)`. The `ptrlen()` primitive above is often a better way to access pointer boundary information, and is less likely to change in future versions.

```

#include "eel.h"
#include "lowlevel.h"

command eel_ptr_example()
{
    char *hello = "Hello world";
    char *p = hello + 6;
    EEL_PTR *ptr = (EEL_PTR *) &p;

    say("Hello's value is %x, size in bytes is %d, base is %x", hello);
    say("P's value is %x, size in bytes is %d, base is %x", p);
    say("P's value is %x, size in bytes is %d, base is %x",
        ptr->value, ptr->size, ptr->base);
    say("P's value is %x", ((int *)&p)[2]);
}

```

The above example shows various ways to display the internal structure of pointers for debugging purposes.

```

char *varptr(int i)
int pointer_to_index(void *)

```

The `varptr()` primitive returns a pointer to any global variable given its index in the name table. The pointer is always a character pointer and should be cast to the correct type before it's used. When `varptr()` is applied to a buffer-specific or window-specific variable, Epsilon checks the `use_default` variable to determine if a pointer to the default or current value should be returned (see page 548). This function doesn't operate with built-in variables—use `get_num_var()` and similar functions for these.

The `pointer_to_index()` primitive does the reverse. It takes a pointer and checks to see if it refers to a global variable. If a global variable is an array or structure, the pointer can point anywhere within. It returns the name table index of the global variable, or 0 if the pointer doesn't point to the contents of any global variable.

```

int vartype(int i)

#define TYPE_BYTE      1      /* an 8-bit unsigned integer */
#define TYPE_CHAR      2      /* a 16-bit unsigned character */
#define TYPE_SHORT     3      /* a 16-bit signed integer */
#define TYPE_INT32      4      /* a 32-bit signed integer */
#define TYPE_INT        5      /* a 64-bit signed integer */
#define TYPE_CARRAY     6      /* character array */
#define TYPE_CPTR       7      /* character pointer */
#define TYPE_POINTER    8      /* contains pointers or spots */
#define TYPE_OTHER      9      /* none of the other categories */
int vartype_class(int i)

```

The `vartype()` primitive returns information on the type of a global variable (or buffer-specific or window-specific variable). It takes the index of the variable in the name table and returns one of the above codes if the variable's type is byte, character, short, in32, integer, character array, or character pointer. It returns `TYPE_POINTER` if the variable is a spot or pointer, or a structure or union containing a spot or pointer. For other types of variables, it returns `TYPE_OTHER`. It returns 0 if the given index doesn't refer to a variable.

The `vartype_class()` subroutine can be more convenient than `vartype()`. It returns 1 if the variable (specified by its name table index) has a numeric type, 2 if it has a string type (`TYPE_CARRAY` or `TYPE_CPTR`), and 0 otherwise.

```
int new_variable(char *name, int type, int vtype, ?int length)
```

The `new_variable()` primitive provides a way to create a new variable without having to load a bytecode file. The first argument specifies the name of the variable. The second argument is a type code of the kind returned by the `name_type()` primitive. The code must be `NT_VAR` for a normal variable, `NT_BUFVAR` for a buffer-specific variable, `NT_WINVAR` for a window-specific variable, or `NT_COLSCHEME` for a color scheme. The third argument is a type code of the kind returned by the `vartype()` primitive described above. This code must be one of the following: `TYPE_BYTE`, `TYPE_CHAR`, `TYPE_SHORT`, `TYPE_INT32`, `TYPE_INT`, or `TYPE_CARRAY`. (This parameter is not used when defining a color scheme.) The last argument is a size, which is used only for `TYPE_CARRAY`. The primitive returns the name table index of the new variable, or -1 if it couldn't create the variable.

10.5.9 Buffer-specific and Window-specific Variables

```
char use_default;
```

Epsilon's buffer-specific variables have a value for each buffer. They change when the current buffer changes. When you create a new buffer, you also automatically create a new copy of each buffer-specific variable. The initial value of each newly created buffer-specific variable is set from special default values Epsilon maintains. These values may be set using the variable `use_default`. When `use_default` is nonzero, referencing any buffer-specific variable accesses its default value, not the value for the current buffer. Otherwise, a value particular to the current buffer applies, as usual.

The normal way to reference a variable's default value is to use the `“.default”` syntax described on page 403, not to set `use_default`.

Window-specific variables have a separate value for each window. When you split a window, the newly created window initially has the same values for all variables as the original window. Each window-specific variable also has a default value, which can be referred to in the same way as buffer-specific variables, via the `“.default”` syntax described on page 403 or by setting the `use_default` variable. Epsilon uses the default value to initialize the first window it creates, during startup, and when it creates pop-up windows.

Only the default values of window- and buffer-specific variables are saved in a state file.

```
copy_buffer_variables(int tobuf, int frombuf)
safe_copy_buffer_variables(int tobuf, int frombuf)
```

The `copy_buffer_variables()` primitive sets all buffer-specific variables in the buffer `tobuf` to their values in the buffer `frombuf`. If `frombuf` is zero, Epsilon resets all buffer-specific variables in the buffer `tobuf` to their default values. The `safe_copy_buffer_variables()` subroutine calls `copy_buffer_variables()`, then clears the values of certain variables that should not be copied between buffers; generally these variables are spot variables that must always refer to positions within their own buffers.

10.5.10 Bytecode Files

```
load_commands(char *file)
load_from_path(char *file)    /* control.e */
int load_eel_from_path(char *file, int flags)
```

The `load_commands()` primitive loads a bytecode file of command, subroutine and variable definitions into Epsilon after the EEL compiler has produced it from the `.e` source file. The primitive changes the name provided so that it has the appropriate `.b` extension, then opens and reads the file. The primitive prints a message and aborts to top-level if it cannot find the file or the file name is invalid.

The subroutine `load_from_path()` searches for a bytecode file using the `lookpath()` primitive (see page 513) and loads it using `load_commands()`.

The `load_eel_from_path()` subroutine searches for an EEL source file with the specified name using `lookpath()`. Then it compiles and loads the file. Bits in the `flags` parameter tell it when to report errors to the user. By default, it doesn't. The 1 bit means complain if the file contained errors, and 2 means also complain if no file by that name was found. The 4 bit has it tell `lookpath()` to search the current directory first. The subroutine returns zero if the file compiled and loaded without errors, one if it wasn't found, and two for other errors.

```
int eel_compile(char *file, int use_fsys, char *flags,
               char *errors, int just_check)
```

The `eel_compile()` primitive lets Epsilon run the EEL compiler without having to invoke a command processor. File specifies the name of a file or buffer. If `use_fsys` is nonzero, it names a file; if `use_fsys` is zero, a buffer. The `flags` parameter may contain any desired command line flags. Compiler messages will go to the buffer named `errors`. Unless errors occur or the `just_check` parameter is nonzero, Epsilon will automatically load the result of the compilation. No bytecode file on disk will be modified. Note that when the compiler includes header files, it will always read them from disk, even if they happen to be in an Epsilon buffer.

The primitive returns 0 on success, 1 if the compilation had fatal errors and did not complete, 2 if the compiler could not be located, or -1 if the user aborted.

```
when_loading()          /* EEL subroutine */
```

Any subroutines with the special name `when_loading()` execute as they are read, and then go away. There may be more than one of these functions defined in a single file. (Note: When the last function defined in an EEL file has been deleted or replaced, Epsilon discards all the constant strings defined in that file. So a file that contains only a `when_loading()` function will lose its constant strings as soon as it exits. If a pointer to such a string must be put in a global variable, use the `strsave()` primitive to make a copy of it. See page 542.)

The `autoload_commands()` primitive described below executes any `when_loading()` functions defined in the file, just as `load_commands()` would. Epsilon never arranges for a `when_loading()` function to be autoloading, and will execute and discard such functions as soon as they're loaded. If you run `autoload_commands()` on a file with `when_loading()` functions, Epsilon will execute them twice: once when it initially sets up the autoloading, and once when it autoloading the file.

```
user char *byte_extension;
user char *state_extension;
```

The extensions used for Epsilon's bytecode files and state files may vary with the operating system. Currently, all operating system versions of Epsilon use ".b" for bytecode files, and ".sta" for state files. The `byte_extension` and `state_extension` primitives hold the appropriate extension names for the particular version of Epsilon.

```
autoload(char *name, char *file, int issubr)
autoload_commands(char *file)
```

Epsilon has a facility to define functions that are not loaded into memory until they are invoked. The `autoload()` primitive takes the name of a function to define, and the name of a bytecode file it can be found in. The file name string may be in a temporary area, because Epsilon makes a copy of it.

The primitive's final parameter should be nonzero to indicate that the autoloading function will be a subroutine, or zero if the function will be a command. (Recall that commands are designed to be invoked directly by the user, and may not take parameters, while subroutines are generally invoked by commands or other subroutines, and may take parameters.) Epsilon enters the command or subroutine in its name table with a special code to indicate that the function is an autoloading function: `NT_AUTOLOAD` for commands, or `NT_AUTOSUBR` for subroutines.

When Epsilon wants to call an autoloading function, it first invokes the EEL subroutine `load_from_path()`, passing it the file name from the `autoload()` call. The standard definition of this function is in the file `control.e`. It searches for the file along the `EPSPATH`, as described on page 12, and then loads the file. The `load_from_path()` subroutine reports an error and aborts the calling function if it cannot find the file.

When `load_from_path()` returns, Epsilon checks to see if the function is now defined as a regular, non-autoloading function. If it is, Epsilon calls it. However, it is not necessarily an error if the function is still undefined. Sometimes a function's work can be done entirely by the `when_loading()` subroutines that are run and immediately discarded as a bytecode file loads.

For example, all the work of the `set-color` command was once done by a `when_loading()` function in the EEL file `color.e`. (In recent versions, it no longer uses autoloading.) Loading the corresponding bytecode file automatically ran this `when_loading()` function, which displayed some windows and let the user choose colors. When the user exited from the command, Epsilon discarded the code for the `when_loading()` function that displayed windows and interpreted keys, and finishes loading the bytecode file. The `set-color` command was still defined as a command that autoloading the `color.b` bytecode file, so the next time the user ran this command, Epsilon loaded the file again.

If the autoloading function was called with parameters, but remains undefined after Epsilon tries to autoloading it, Epsilon aborts the calling function with an error message. Functions that use the above technique to load temporarily may not take parameters.

Like `load_commands()`, the primitive `autoload_commands()` takes the name of a compiled EEL bytecode file as a parameter. It loads any variables or bindings contained in the file, just like `load_commands()`. But instead of loading the functions in the file, this primitive generates an autoloading request for each function in the file. Whenever any EEL function tries to call a function in the file, Epsilon will load the entire file.

10.5.11 Starting and Finishing

```
do_save_state(char *file)
int save_state(char *file)
```

The `do_save_state()` subroutine writes the current state to the specified file. It aborts with an error message if it encounters a problem. It uses the `save_state()` primitive to actually write the state. The

primitive returns 0 if the information was written successfully, or an error code if there was a problem (as with `file_write()`). Both change the extension to “.sta” before using the supplied name.

The state includes all commands, subroutines, keyboard macros, and variables. It does not include buffers or windows. Since a state file can only be read while Epsilon is starting (when there are no buffers or windows), only the default value of each buffer-specific or window-specific variable is saved in a state file.

Pointer variables will have a value of zero when the state file is loaded again. Epsilon does not save the object that is pointed to. Spot variables and structures or unions containing pointers or spots are also zeroed, but other types of variables are retrieved unchanged (but see the description of the `zeroed` keyword on page 424).

```
short argc;
char *argv[ ];
```

When Epsilon starts, it examines the arguments on its command line, and modifies its behavior if it recognizes certain special flags. But first it adds in the contents of the configuration variable `EPSILON`, if this exists, putting this before any actual command line parameters.

Epsilon looks for certain special flags, interprets them and removes them from the command line. It then passes the remainder of the command line to the EEL startup code in `cmdline.e`. That code interprets any remaining flags and files on the command line. You can add new flags to Epsilon by modifying `cmdline.e`. See page 13 for the meaning of each of Epsilon’s flags.

Epsilon interprets and removes these flags from the command line:

-k	Keyboard options	-m	Memory control
-s	Load from state file	-w	Directory options
-b	Load from bytecode file	-v	Video options

Some of these settings are visible to an EEL program through variables. See the `load-from-state` variable for the `-b` flag, the `state_file` variable for the `-s` flag, the `want-cols` and `want-lines` variables for the `-vc` and `-vl` flags, and the `directory-flags` variable for the `-w` flag.

All other flags, as well as any specified files, are interpreted by the EEL functions in `cmdline.e`. They read the command line from the `argc` and `argv` variables, already broken down into words. The `argc` variable contains the number of words in the command line. The `argv` variable contains the words themselves. The first word on the command line, `argv[0]`, is always the name of Epsilon’s executable file, so that if `argc` is 2, there was one argument and it is in `argv[1]`.

```
char *original_argv(int n)
```

The `original_argv()` primitive lets EEL code access Epsilon’s original command line arguments, including those interpreted internally and not passed along to EEL code in the `argv` array, and excluding those added by any `EPSILON` configuration variable. Calling `original_argv(1)` returns the first command line argument, `original_argv(2)` the second, and so forth. A null return value indicates there are no more arguments. Calling `original_argv(0)` returns the full path of Epsilon’s executable.

```
when_restoring()      /* cmdline.e */
early_init()          /* cmdline.e */
middle_init()         /* cmdline.e */
start_up()            /* cmdline.e */
user char *version;
apply_defaults()
```

Epsilon calls the EEL subroutine `when_restoring()` if it exists after loading a state file. Unlike `when_loading()`, this subroutine is not removed after it executes. The standard version of `when_restoring()` sets up variables and modes, and interprets the command line. It calls several EEL subroutines at various points in the process. Each does nothing by default, but you can conveniently customize Epsilon by redefining them. (See page 544 to make sure your extension doesn't interfere with other extensions.)

The `when_restoring()` function calls `early_init()` just before interpreting flags, and `middle_init()` just after. It then loads files (from the command line, or a saved session), displays Epsilon's version number, and calls the `start_up()` subroutine. (The `version` variable contains a string with the current version of Epsilon, such as "9.0".) Finally, Epsilon executes any `-l` and `-r` switches.

The `when_restoring()` subroutine calls the `apply_defaults()` primitive before it calls `early_init()`. This primitive sets the values of window-specific and buffer-specific variables in the current buffer and window to their default values.

```
char state_file[ ];
user char load_from_state;
```

The `state_file` primitive contains the name of the state file Epsilon was loaded from, or "" if it was loaded only using bytecode files with the `-b` flag. The `load_from_state` variable will be set to 1 if Epsilon loaded its functions from a state file at startup, or 0 if it loaded only from bytecode files.

```
after_loading()
```

After Epsilon calls the `when_restoring()` subroutine, it finishes its internal initialization by checking for the existence of certain variables and functions that must be defined if Epsilon is to run. Until this is done, Epsilon can't perform a variety of operations such as getting a key from the keyboard, displaying buffers, and searching. The `after_loading()` primitive tells Epsilon to finish initializing now. The variables and functions listed in figure 10.2 must be defined when you call `after_loading()`.

```
finish_up()
user char leave_blank;
```

When Epsilon is about to exit, it calls the subroutine `finish_up()`, if it exists. (See page 544 to make sure your extension doesn't interfere with other extensions that may also define `finish_up()`.) Epsilon normally redisplay each mode line one last time just before exiting, so any buffers that it saved just before exiting will not still be marked unsaved on the screen. However, if the `leave_blank` primitive is nonzero, it skips this step.

10.5.12 EEL Debugging and Profiling

```
int name_debug(int index)
set_name_debug(int index, int flag)
```

Every command or subroutine in Epsilon's name table has an associated debug flag. If the debug flag of a command or subroutine is nonzero, Epsilon will start up the EEL debugger when the function is called, allowing you to step through the function line by line. See page 169. The `name_debug()` primitive returns the debug flag for an item, and the `set_name_debug()` primitive sets it.

```
when_idle()
when_displaying()
when_repeating()
getkey()
on_modify()
prepare_windows()
build_mode()
fix_cursor()
load_from_path()
color_class standard_color;
color_class standard_mono;
user int see_delay;
user short beep_duration;
user short beep_frequency;
user char mention_delay;
char _display_characters[ ];
user buffer int undo_size;
buffer short *mode_keys;
user buffer short tab_size;
user buffer short case_fold;
buffer char *_display_class;
char *_echo_display_class;
user window int display_column;
window char _highlight_control;
window char _window_flags;
char use_process_current_directory;
```

Figure 10.2: Variables and functions that must be defined.

```

start_profiling()
stop_profiling()
char *get_profile()

```

Epsilon can generate an execution profile of a section of EEL code. A profile is a tool to determine which parts of a program are taking the most time. The `start_profiling()` primitive begins storing profiling information internally. Profiling continues until Epsilon runs out of space, or you call the `stop_profiling()` primitive, which stops storing the information. Many times each second, Epsilon saves away information describing the location in the source file of the EEL code it is executing, if you've turned profiling on. You can use this to see where a command is spending its time, so that you can center your efforts to speed the command up there.

Once you stop the profiling with the `stop_profiling()` primitive, you can retrieve the profiling information with the `get_profile()` primitive. Each call returns one line of the stored profile information, and the function returns a null pointer when all the information has been retrieved. Each line contains the name of an EEL source file and a line number within the file, separated by a space. See the `profile` command for a more convenient way to use these primitives. Functions that you've compiled with the EEL compiler's `-s` flag will not appear in the profile.

10.5.13 Help Subroutines

```

int name_help(int index)
set_name_help(int index, int offset)
get_doc()           /* help.e */

```

Every item in Epsilon's name table has an associated help file offset. The help offset contains the position in Epsilon's help file "edoc" where information on an item is stored. Epsilon uses it to provide quick access to help file items. It is initially -1, and may be set with the `set_name_help()` primitive and examined with the `name_help()` primitive. (The Windows version of Epsilon normally uses a standard Windows help file to display help, so it doesn't use these help file offsets.)

When an EEL function wants to look up information in the help file, it calls the EEL subroutine `get_doc()`. This function loads the help file into the buffer "-edoc" if it hasn't been loaded before.

Epsilon's help file "edoc" uses a simple format that makes it easy to add new entries for your own commands. Each command's description begins with a line consisting of a tilde (~), the command or variable's name, a <Tab>, and the command's one-line description (or, for a variable, some type information). Following lines (until the next line that starts with ~, or the end of the file) constitute the command's full description. The entries can occur in any order; they don't have to be listed alphabetically.

An entry can contain a cross-reference link to another entry in the file; these consist of the name of the command or variable being cross-referenced, bracketed by two control characters. Put a ^A character before the name of the command or variable, and a ^B character after. Also see the description of the `view_linked_buf()` subroutine on page 464.

```

help_on_command(int ind)           /* help.e */
help_on_current()                 /* help.e */

```

The `help_on_command()` subroutine provides help on a particular command. It takes the name table index of the command to provide help on.

The `help_on_current()` subroutine displays help on the currently-running command. It uses the `last_index` variable to determine the current command.

```
show_binding(char *fmt, char *cmd) /* help.e */
```

The `show_binding()` subroutine displays the message `fmt` using the `say()` primitive. The `fmt` must contain the `%s` sequence (and no other `%` sequences). Epsilon will replace the `%s` with the binding of the command `cmd`. For example,

```
show_binding("Type %s to continue", "exit-level");
```

displays “Type Ctrl-x Ctrl-z to continue” with Epsilon’s normal bindings.

10.6 Input Primitives

10.6.1 Keys

```
wait_for_key()
user int key;
user int full_key;
int generic_key(int k)
when_idle(int times) /* EEL subroutine */
add_buffer_when_idle(int buf, int (*func)())
delete_buffer_when_idle(int buf, int (*func)())
when_repeating() /* EEL subroutine */
int is_key_repeating()
```

The `wait_for_key()` primitive advances to the next key, waiting for one if necessary. The variable `key` stores the last key obtained from `wait_for_key()`.

Some key combinations have both a generic and specific interpretation. For instance, the `<Plus>` key on a numeric keypad has a generic code identical to the `<Plus>` key on the main keyboard, but a unique specific code. Epsilon sets the `full_key` variable to the specific key code for that key, and `key` to the generic code for the key (in this case, “+”). For keys with only one interpretation, Epsilon sets `key` and `full_key` the same.

The `generic_key()` primitive returns the generic version of the specified key. For keys with only one interpretation, it returns the original key code.

When you call `wait_for_key()`, it first checks to see if the `ungot_key` variable has a key (see below) and uses that if it does. If not, and a keyboard macro is active, `wait_for_key()` returns the next character from the macro. (The primitive also keeps track of repeat counts for macros.) If there is no key in `ungot_key` and no macro is active, the primitive checks to see if you have already typed another key and returns it if you have. If not, the primitive waits until you type a key (or a mouse action or other event occurs—Epsilon treats all of these as keys).

When a concurrent process is outputting text into an Epsilon buffer, it only appears there during a call to `wait_for_key()`. Epsilon handles the processing of other concurrent events like FTP transfers during this time as well.

While Epsilon is waiting for a key, it calls the `when_idle()` subroutine. The default version of this function does idle-time code coloring and displays any defined idle-time message in the echo area (see the `show-when-idle` variable), among other things. The `when_idle()` subroutine receives a parameter that indicates the number of times the subroutine has been called since Epsilon began waiting for a key. Every time Epsilon gets a key (or other event), it resets this count to zero.

The `when_idle()` subroutine should return a timeout code in hundredths of a second. Epsilon will not call the subroutine again until the specified time has elapsed, or another key arrives. If it doesn't need Epsilon to call it for one second, for example, it can return 100. If it wants Epsilon to call it again as soon as possible (assuming Epsilon remains idle), it can return 0. If the subroutine has completed all its work and doesn't need to be called again until after the next keystroke or mouse event, it can return -1. Epsilon will then go idle waiting for the next event. (The return value is only advisory; Epsilon may call `when_idle()` more frequently or less frequently than it requests.)

A mode may wish to provide additional functions that run during idle time, beyond those the `when_idle()` subroutine performs itself. The `add_buffer_when_idle()` subroutine registers a function `func` so that it will be called during idle-time processing whenever `buf` is the current buffer. The `delete_buffer_when_idle()` subroutine removes the specified function from that buffer's list of buffer-specific idle-time functions. (It does nothing if the function was not on the list.) A buffer-specific when-idle function takes a parameter `times` and must return a result in the same fashion as the `when_idle()` function itself.

To add an idle-time task not associated with any specific buffer, or one that runs even if a given buffer isn't the current one, define a function with a name that starts with `do_when_idle_`. Epsilon will call it whenever it's idle. It must take a parameter `times` and return a result just like the `when_idle()` function.

When you hold down a key to make it repeat, Epsilon does not call the `when_idle()` subroutine. Instead, it calls the `when_repeating()` subroutine. Again, this varies by environment: under some operating systems, Epsilon cannot distinguish between repeated key presses and holding down a key to make it repeat. If this is the case, Epsilon won't call the function.

You can add your own logic for when a key repeats by defining a function with a name that starts with `do_when_repeating_`. The `when_repeating()` subroutine will call it whenever it runs. It must take no parameters and return no result.

The `is_key_repeating()` primitive returns nonzero if the user is currently holding down a key causing it to repeat. Epsilon can't detect this in all environments, so the primitive always returns 0 in that case.

```
int getkey()          /* control.e */
```

Instead of calling `wait_for_key()` directly, EEL commands should call the EEL subroutine `getkey()` (defined in `control.e`), to allow certain actions that are written in EEL code to take effect on each character. For example, the standard version of `getkey()` saves each new character in a macro, if you're defining one. It checks the EEL variable `_len_def_mac`, which contains the length of the macro being defined plus one, or zero if you're not defining a macro. For convenience, `getkey()` also returns the new key. The `getkey()` subroutine calls `wait_for_key()`. (If you want to add functions to `getkey()`, see page 544 to make sure your extension doesn't interfere with other extensions that may also add to `getkey()`.)

```
int char_avail()
int in_macro(?int ignore_suspended)
```

The `char_avail()` primitive returns 0 if `wait_for_key()` would have to wait if it were called, and 1 otherwise. That is, it returns nonzero if and only if a key is available from `ungot_key`, a keyboard macro, or the keyboard.

The `in_macro()` primitive returns 1 if a keyboard macro is running or has been suspended, 0 otherwise. If its optional `ignore_suspended` parameter is nonzero, it counts a suspended macro as if no macro were running. While processing the last key of a keyboard macro, `in_macro()` will return 0, because Epsilon has already discarded the keyboard macro by that time. Check the `key-from-macro` variable instead to see if the key currently being handled came from a macro.

There are some textual macros defined in `eel.h` which help in forming the codes for keys in an EEL function. The codes for normal ASCII keys are their ASCII codes, so the code for the key 'a' is 'a'; the same goes for Unicode characters. The `ALT()` macro makes these normal keys into their Alt forms, so the code for Alt-a is `ALT('a')`. The `CTRL()` macro changes a character into the corresponding control character, so `CTRL('h')` or `CTRL('H')` both represent the Ctrl-h key. Both `CTRL(ALT('q'))` and `ALT(CTRL('q'))` stand for the Ctrl-Alt-q key.

The remaining key codes represent those keys that don't correspond to any possible buffer character, plus various key codes that represent other kinds of input events, such as mouse activity.

The `FKEY()` macro represents the function keys. `FKEY(1)` and `FKEY(12)` are F1 and F12, respectively. Note that this macro takes a number, not a character.

Refer to the cursor pad keys using the macros `KEYINSERT`, `KEYEND`, `KEYDOWN`, `KEYPGDN`, `KEYLEFT`, `KEYRIGHT`, `KEYHOME`, `KEYUP`, `KEYPGUP`, and `KEYDELETE`. You can refer to the numeric keypad keys with the `NUMDIGIT()` macro: `NUMDIGIT(0)` is N-0, and `NUMDIGIT(9)` is N-9. `NUMDOT` is the numeric keypad period, and `NUMENTER` is the `<Enter>` or `<Return>` key on the numeric keypad (normally mapped to Ctrl-m).

The codes for the remaining keys are `GREYPLUS`, `GREYMINUS`, `GREYSTAR`, `GREYSLASH`, `GREYEQUAL`, and `GREYHELP` for the +, -, *, /, =, and Help keys on the numeric keypad (not every keyboard has all these keys), and `GREYENTER`, `GREYBACK`, `GREYTAB`, `GREYESC`, and `SPACEBAR` for the `<Enter>`, `<Backspace>`, `<Tab>`, `<Esc>`, and `<Spacebar>` keys, respectively.

```
#define NUMSHIFT(c)      ((c) | KEY_SHIFT)
#define NUMCTRL(c)       ((c) | KEY_CTRL)
#define NUMALT(c)        ((c) | KEY_ALT)
#define KEY_PLAIN(c)     ((c) & ~ (KEY_SHIFT | KEY_CTRL | KEY_ALT))
```

The `NUMSHIFT()`, `NUMCTRL()`, and `NUMALT()` macros make shifted, control, and alt versions of keys, respectively, by turning on the bit in a key code for each of these properties: `KEY_SHIFT`, `KEY_CTRL`, and `KEY_ALT`. For example, `NUMCTRL(NUMDIGIT(3))` is Ctrl-n-`<PgDn>`, and `NUMALT(KEYDELETE)` is A-``. The `KEY_PLAIN()` macro strips away these bits.

In this version, `NUMALT()` and `ALT()` are the same, and `NUMCTRL()` and `CTRL()` only differ on certain low-numbered characters: the former always turns on the `KEY_CTRL` bit, while the latter also generates ASCII control codes when given suitable ASCII characters.

```
int make_alt(int k)      /* control.e */
int make_ctrl(int k)     /* control.e */
```

The `make_alt()` subroutine defined in `control.e` will return an Alt version of any key. The `make_ctrl()` subroutine is similar, but makes a key into its Control version. These may be used instead of the `ALT()` and `CTRL()` macros.

Use the `IS_CTRL_KEY()` macro to determine if a given key is a control key of some kind. Its value is nonzero if the key is an ASCII Control character, a function key with Control held down, or any other Control key. It understands all types of keys. The macro `IS_ALT_KEY()` is similar; its value is nonzero if the given key was generated when holding down the Alt key.

A macro command recorded using the notation `<!find-file>` uses the bit flag `CMD_INDEX_KEY`. In this case the value of key is not a true key, but rather the name table index of the specified command. See page 172 for more information.

```
user int ungot_key;
```

If the `ungot_key` variable is set to some value other than its usual value of `-1`, that number is placed in `key` and `full_key` as the new key when `wait_for_key()` is called next, and `ungot_key` is set to `-1` again. You can use this to make a command that reads keys itself, then exits and runs the key again when you press an unrecognized key. The statement `ungot_key = key;` accomplishes this.

```
show_char(char *str, int key, ?int style)
```

The `show_char()` primitive converts a key code to its printed representation, described on page 172. For example, the code produced by function key 3 generates the string `F-3`. The string is *appended* to the character array `str`.

If `show_char()`'s optional third parameter is present, and nonzero, this primitive will use a longer, more readable printed representation. For example, rather than `C-A-S` or `,` or `S-F-10`, `show_char()` will return `Ctrl-Alt-s` or `<Comma>` or `Shift-F10`. (Epsilon can only parse the former style, in Epsilon command files and in all other commands that use the `get_keycode()` primitive below.)

This function always represents non-Latin1 Unicode characters (those in the range 256–65535) with their character names, like `<GREEK SMALL LETTER GAMMA>`. With a style of 3, Epsilon does this for Latin 1 characters (those in the range 32–255) too, thus representing all Unicode characters by name. With a style of 2, Epsilon uses Unicode character names for non-ASCII Latin 1 characters (those in the range 128–255) but represents printable ASCII characters like `"J"` as-is. With a style of 1, it represents all Latin 1 characters (in the range 32–255) as-is, using Unicode character names only for characters over 255.

See the `%k` sequence used by `sprintf()` and other functions, described on page 480, for a more convenient way to translate keys or characters to text. It always uses style 1. Also see the `unicode_character_name()` primitive.

```
#define key_t          int
key_t *get_keycode()
int key_value(char *s, ?char **after)
stuff_macro(key_t *mac, int oneline)
```

The `get_keycode()` primitive is used to translate a sequence of key names such as `"C-xC-A-f"` into the equivalent key codes. It moves past a quoted sequence of key names in the buffer and returns an array of ints with the key codes. The same array is used each time the function is called. The first entry of the array contains the number of array entries. The primitive returns null if the string had an invalid key name.

The `key_t` macro represents the type of a key. It's the same as an `int` in this versions, but not in older versions. Writing `key_t` instead of `int`, and `#including` the compatibility header file `oldkeys.h` helps make EEL code compatible with older versions.

The `key_value()` primitive also converts key names into key codes, but it gets the key name from a string, not a buffer, and returns a single key code at a time. It tries to interpret `s` as a key name, and returns its value. If the optional pointer `after` is non-null, it must point to a character pointer. Epsilon sets `*after` to the position in the string after the key name. When `s` contains an invalid key name, `key_value()` returns `-1` and sets `*after` (if `after` is non-null) to `s`.

The `stuff_macro()` subroutine inserts a sequence of key names into the current buffer in a format that `get_keycode()` can read, surrounding the key names with `"` characters. The list of keys is specified by an array of ints in the same format `get_keycode()` uses: the first value contains the total number of array entries. If `oneline` is nonzero, the subroutine represents line breaks with `\n` so that the text stays on one line.

```
user char key_type;
user short key_code;
```

When `wait_for_key()` returns a key that comes directly from the keyboard, it also sets the primitive variables `key_type` and `key_code`. These let EEL programs distinguish between keys that translate to the same Epsilon key code, for certain special applications. The `wait_for_key()` primitive doesn't change either variable when the key comes from `ungot_key`.

The `key_code` variable contains the sixteen-bit BIOS-style encoding for the key that Epsilon received from the operating system, if available. Its ASCII code is in the low eight bits and its scan code is in the high eight bits.

The `key_type` variable has one of the following values, defined in `codes.h`. If `KT_NONASCII` or `KT_NONASCII_EXT`, the key was a special key without an ASCII translation, such as a function key. Such keys are of type `KT_NONASCII_EXT` if they're one of the keys on an extended keyboard that are synonyms to multikey sequences on the old keyboard, such as the keys on the extended keyboard's cursor pad.

A key type of `KT_ACCENT_SEQ` indicates a multikey sequence that the operating system or a resident program has translated as a single key, such as an `ê`. Key type `KT_ACCENT` generally means the operating system translated a single key to a graphics character or foreign language character. Key type `KT_NORMAL` represents any other key. Most keys have a key type of `KT_NORMAL`.

A key type of `KT_MACRO` means the key came from a macro. A macro key recorded with the `EXTEND_SEL_KEY` bit flag returns a key type of `KT_EXTEND_SEL` instead, but these extend codes are not used in current versions of Epsilon. In either case, the `key_code` variable is set to zero in this case.

In many environments, the `key_code` variable is always zero, and `key_type` is either `KT_NORMAL`, `KT_MACRO`, or `KT_EXTEND_SEL`.

10.6.2 The Mouse

When a mouse event occurs, such as a button press or a mouse movement, Epsilon enqueues the information in the same data structure it uses for keyboard events. A call to `wait_for_key()` retrieves the next item from the queue—either a keystroke or a mouse event. Normally an EEL program calls the `getkey()` subroutine instead of `wait_for_key()`. See page 555.

```
user short catch_mouse;
```

The `catch_mouse` primitive controls whether Epsilon will queue up any mouse events. Setting it to zero causes Epsilon to ignore the mouse. A nonzero value makes Epsilon queue up mouse events. If your system has no mouse, setting `catch_mouse` has no effect.

```
user short mouse_mask;
user short mouse_x, mouse_y;
user short mouse_screen;
user int double_click_time;
```

You can control which mouse events Epsilon dequeues, and which it ignores, by using the `mouse_mask` primitive. The following values, defined in `codes.h`, control this:

```
#define MASK_MOVE      0x01
#define MASK_LEFT_DN   0x02
#define MASK_LEFT_UP   0x04
#define MASK_RIGHT_DN  0x08
#define MASK_RIGHT_UP  0x10
#define MASK_CENTER_DN 0x20
```

```
#define MASK_CENTER_UP 0x40
#define MASK_ALL      0x7f
#define MASK_BUTTONS  (MASK_ALL - MASK_MOVE)
#define MASK_DN       // ... see eel.h
#define MASK_UP       // ... see eel.h
```

For example, the following EEL code would cause Epsilon to pay attention to the left mouse button and mouse movement, but ignore everything else:

```
mouse_mask = MASK_MOVE | MASK_LEFT_DN | MASK_LEFT_UP;
```

When Epsilon dequeues a mouse event with `wait_for_key()`, it sets the values of `mouse_x` and `mouse_y` to the screen coordinates associated with that mouse event. Setting them moves the mouse cursor. The upper left corner has coordinate (0, 0).

When dequeuing a mouse event, `wait_for_key()` returns one of the following “keys” (defined in `codes.h`):

```
MOUSE_LEFT_DN    MOUSE_LEFT_UP    MOUSE_DBL_LEFT
MOUSE_CENTER_DN  MOUSE_CENTER_UP  MOUSE_DBL_CENTER
MOUSE_RIGHT_DN   MOUSE_RIGHT_UP   MOUSE_DBL_RIGHT
MOUSE_MOVE
```

Dequeuing a mouse event also sets the `mouse_screen` variable to indicate which screen its coordinates refer to. Screen coordinates are relative to the specified screen. Ordinary Epsilon windows are on the main screen, screen 0. When Epsilon creates a dialog box containing Epsilon windows, each Epsilon window receives its own screen number. For example, if you type `Ctrl-x Ctrl-f ?`, Epsilon displays a dialog box with two screens, usually numbered 1 and 2. If you click on the ninth line of the second screen, Epsilon returns the key `MOUSE_LEFT_DN`, sets `mouse_y` to 8 (counting from zero), and sets `mouse_screen` to 2.

The `double_click_time` primitive specifies how long a delay to allow for double-clicks (in hundredths of a second). If two consecutive `MOUSE_LEFT_DN` events occur within the allotted time, then Epsilon enqueues a `MOUSE_DBL_LEFT` event in place of the second `MOUSE_LEFT_DN` event. The corresponding thing happens for right clicks and center clicks as well. Epsilon for Windows ignores this variable and uses standard Windows settings to determine double-clicks.

```
#define IS_WIN_KEY(k)      // ... omitted
#define IS_MOUSE_KEY(k)   // ... omitted
#define IS_TRUE_KEY(k)    // ... omitted
#define IS_EXT_ASCII_KEY(k) // ... omitted
#define IS_WIN_PASSIVE_KEY(k) // ... omitted
#define IS_MOUSE_LEFT(k)  // ... omitted
#define IS_MOUSE_RIGHT(k) // ... omitted
#define IS_MOUSE_CENTER(k) // ... omitted
#define IS_MOUSE_SINGLE(k) // ... omitted
#define IS_MOUSE_DOUBLE(k) // ... omitted
#define IS_MOUSE_DOWN(k)  // ... omitted
#define IS_MOUSE_UP(k)    // ... omitted
```

The `IS_MOUSE_KEY()` macro returns a nonzero value if the given key code indicates a mouse event. The `IS_TRUE_KEY()` macro returns a nonzero value if the given key code indicates a keyboard key. The

IS_EXT_ASCII_KEY() macro returns a nonzero value if the given key code represents a character that can appear in a buffer (rather than a function key or cursor key). The IS_WIN_KEY() macro returns a nonzero value if the given key code indicates a window event like a menu selection, pressing a button on a dialog, or getting the focus, while IS_WIN_PASSIVE_KEY() returns nonzero if the given key represents an incidental event: losing or gaining focus, losing the selection, or the mouse entering or leaving Epsilon's window.

The IS_MOUSE_LEFT(), IS_MOUSE_RIGHT(), and IS_MOUSE_CENTER() macros return nonzero if a particular key code represents either a single or a double click of the indicated button. The IS_MOUSE_SINGLE() and IS_MOUSE_DOUBLE() macros return nonzero if the given key code represents a single-click or double-click, respectively, of any mouse button. The IS_MOUSE_DOWN() macro returns nonzero if the key code represents the pressing of any mouse button (either a single-click or a double-click). Finally, the IS_MOUSE_UP() macro tells if a particular key code represents the release of any mouse button.

```
user short mouse_pixel_x, mouse_pixel_y;
int y_pixels_per_char()
int x_pixels_per_char()
clip_mouse() /* mouse.e subr. */
```

On most systems, Epsilon can provide the mouse position with finer resolution than simply which character it is on. The mouse_pixel_x and mouse_pixel_y variables contain the mouse position in the most accurate form Epsilon provides. Setting the pixel variables moves the mouse cursor and resets the mouse_x and mouse_y variables to match. Similarly, setting mouse_x or mouse_y resets the corresponding pixel variable.

EEL subroutines should not assume any particular scaling between the screen character coordinates provided by mouse_x and mouse_y and these "pixel" variables. The scaling varies with the screen display mode or selected font. As with the character coordinates, the upper left corner has pixel coordinate (0, 0). The y_pixels_per_char() and x_pixels_per_char() primitives report the current scaling between pixels and characters. For example, mouse_x usually equals the quantity mouse_pixel_x / x_pixels_per_char(), rounded down to an integer.

The mouse_x variable can range from -1 to screen_cols, while the valid screen columns range from 0 to (screen_cols - 1). Epsilon uses the additional values to indicate that the user has tried to move the mouse cursor off the screen, in environments which can detect this (currently, no supported environments can). The mouse_pixel_x variable, on the other hand, ranges from 0 to screen_cols * x_pixels_per_char(). The highest and lowest values of mouse_pixel_x correspond to the highest and lowest values of mouse_x, while other values obey the relation outlined in the previous paragraph. The mouse_y and mouse_pixel_y variables work in the same way.

The clip_mouse() subroutine alters the mouse_x and mouse_y variables so that they refer to a valid screen column, if they currently range off the screen.

```
user short mouse_shift;
short shift_pressed()
#define KB_ALT_DN      0x08    // Some Alt key
#define KB_CTRL_DN     0x04    // Some Ctrl key
#define KB_LSHIFT_DN  0x02    // Left shift key
#define KB_RSHIFT_DN  0x01    // Right shift key
#define KB_SHIFT_DN   (KB_LSHIFT_DN | KB_RSHIFT_DN)
                        // Either shift key

int was_key_shifted()
```

When Epsilon dequeues a mouse event with wait_for_key(), it also sets the mouse_shift variable to indicate which shift keys were depressed at the time the mouse event was enqueued. The

`shift_pressed()` primitive returns the same codes, but indicates which shift keys are depressed at the moment you call it.

The `was_key_shifted()` subroutine tells if the user held down Shift when pressing the current key. Some keys produce the same key code with or without shift.

Unlike the `shift_pressed()` primitive, which reports on the current state of the Shift key, this one works with keyboard macros by returning the state of the Shift key at the time the key was originally pressed. A subroutine must call `was_key_shifted()` at the time the macro is recorded for the Shift state to be recorded in the macro. Macros defined by a command file can use an E- prefix to indicate this.

```
short mouse_buttons()
int mouse_pressed()
get_movement_or_release() /* menu.e */
```

The `mouse_buttons()` primitive returns the number of buttons on the mouse. A value of zero means that Epsilon could not find a mouse on the system.

The `mouse_pressed()` primitive returns a nonzero value if and only if some button on the mouse has gone down but has not yet gone up. The subroutine `get_movement_or_release()` uses this function. It delays until the mouse moves or all its buttons have been released.

Mouse Cursors

```
user short mouse_display;
user short mouse_auto_on; /* default = 1 */
user short mouse_auto_off; /* default = 1 */
```

The `mouse_display` primitive controls whether or not Epsilon displays the mouse cursor. Set it to zero to turn the mouse cursor off, and to a nonzero value to turn the mouse cursor on. Turning off the mouse cursor does not cause Epsilon to stop queuing up mouse events—to do that, use `catch_mouse`.

Epsilon automatically turns on the mouse cursor when it detects mouse motion, if the `mouse_auto_on` primitive has a nonzero value. Epsilon automatically turns off the mouse when you start to type on the keyboard, if the `mouse_auto_off` primitive has a nonzero value. Neither of these actions affect the status of queuing up mouse events. When Epsilon automatically turns on the mouse cursor, it sets `mouse_display` to 2.

```
typedef short MOUSE_CURSOR;
MOUSE_CURSOR mouse_cursor;

MOUSE_CURSOR std_pointer      = CURSOR_IBEAM;
MOUSE_CURSOR left_pointer     = CURSOR_ARROW;
MOUSE_CURSOR drag_and_drop_pointer = CURSOR_ARROW;
MOUSE_CURSOR right_pointer    = CURSOR_RIGHTARR;
MOUSE_CURSOR horiz_arrows     = CURSOR_SIZEWE;
MOUSE_CURSOR vert_arrows      = CURSOR_SIZENS;
MOUSE_CURSOR move_arrows      = CURSOR_ARROW;
MOUSE_CURSOR pan_neutral      = CURSOR_PAN;
MOUSE_CURSOR pan_up           = CURSOR_PAN_UP;
MOUSE_CURSOR pan_down         = CURSOR_PAN_DOWN;
```

Setting the `mouse_cursor` primitive selects a different mouse cursor in Windows and X11 versions of Epsilon. It holds a code that specifies the type of cursor to use. Epsilon regularly sets this variable from an appropriate one of the mouse cursor variables shown above (for instance, from `pan_up` while panning up using the mouse's middle button). By default, Epsilon sets `mouse_cursor` to `std_pointer`, which produces an I-beam cursor.

The codes in all these variables select one of the standard cursor types, according to the following table, which lists the stock cursor codes defined in `codes.h`:

<code>CURSOR_ARROW</code>	Standard arrow
<code>CURSOR_IBEAM</code>	Text I-beam
<code>CURSOR_WAIT</code>	Hourglass
<code>CURSOR_CROSS</code>	Crosshair
<code>CURSOR_UPARROW</code>	Arrow pointing up
<code>CURSOR_SIZE</code>	Resize
<code>CURSOR_ICON</code>	Empty icon
<code>CURSOR_SIZENWSE</code>	Double-headed arrow pointing northwest and southeast
<code>CURSOR_SIZENESW</code>	Double-headed arrow pointing northeast and southwest
<code>CURSOR_SIZEWE</code>	Double-headed arrow pointing east and west
<code>CURSOR_SIZENS</code>	Double-headed arrow pointing north and south
<code>CURSOR_PAN</code>	Neutral cursor for wheeled mouse panning
<code>CURSOR_PAN_UP</code>	Wheeled mouse cursor when panning up
<code>CURSOR_PAN_DOWN</code>	Wheeled mouse cursor when panning down
<code>CURSOR_HAND</code>	A hand
<code>CURSOR_APPSTART</code>	arrow with hourglass
<code>CURSOR_HELP</code>	arrow with question mark
<code>CURSOR_RIGHTARR</code>	arrow pointing up and right

X11 defines many other mouse cursors. You can use the number of any X11 mouse cursor code by setting Epsilon's mouse cursor variable to its X11 cursor font glyph number plus 50. (For example, X11's familiar "fat X" default mouse cursor has code 0, so `mouse_cursor = 50` produces it, while X11's less familiar "trek" mouse cursor has code 142, so `mouse_cursor = 192` produces it.)

Mouse Subroutines

```
window int (*mouse_handler)();
allow_mouse_switching(int nwin)    // mouse.e subr.
buffer char mouse_dbl_selects;
char run_by_mouse;
char show_mouse_choices;
```

The `mouse.e` and `menu.e` files define the commands and functions normally bound to the mouse buttons. The functions that handle button clicks examine the window-specific function pointer `mouse_handler` so that you can easily provide special functions for clicks in a particular window. By default, the variable contains 0 in each window, so that Epsilon does no special processing. Set the variable to point to a function, and Epsilon will call it whenever the user pushes a mouse button and the mouse cursor is over the indicated window. The function receives one parameter, the window handle of the specified window. It can return nonzero to prevent the normal functioning of the button, or zero to let the function proceed.

The `allow_mouse_switching()` subroutine is a `mouse_handler` function. Normally, when a pop-up window is on the screen, Epsilon doesn't let the user simply switch to another window. Depending on the context, Epsilon either removes the pop-up window and then switches to the new window, or signals an error and remains in the pop-up window. If you set the `mouse_handler` variable in a particular window to the `allow_mouse_switching()` subroutine, Epsilon will permit switching to that window if the user clicks in it, without deleting any pop-up window.

The buffer-specific `mouse_dbl_selects` variable controls what double-clicking with a mouse button does. By default the variable is zero, and double-clicking selects words. If the variable is nonzero, Epsilon instead runs the command bound to the `<Newline>` key.

The `run_by_mouse` variable is normally zero. Epsilon sets it to one while it runs a command that was selected via a pull-down menu or using the tool bar. Commands can use this variable to behave differently in this case. For example, the subroutine that provides completion automatically produces a list of choices to choose from, when run via the mouse. It does this if the `MUST_MATCH` flag (see page 568) indicates that the user must always pick one of the choices (instead of typing in a different selection), or if the `show-mouse-choices` variable is nonzero.

The Scroll Bar

```
user window int display_scroll_bar;
int scroll_bar_line()
```

The built-in variable `display_scroll_bar` controls whether or not the current window's right border contains a scroll bar. Set it to zero to turn off the scroll bar, or to any positive number to display the bar. If a window has no right border, or has room for fewer than two lines of text, Epsilon won't display a scroll bar. Although the EEL functions that come with Epsilon don't support clicking on a scroll bar on the left border of a window, Epsilon will display one if `display_scroll_bar` is negative. Any positive value produces the usual right-border scroll bar. (This variable, and the following primitive, have no effect in Epsilon for Windows, which handles scrolling internally.)

The `scroll_bar_line()` primitive returns the position of the scroll box diamond on the scroll bar. A value of one indicates the line just below the arrow at the top of the scroll bar. Epsilon always positions this arrow adjacent to the first line of text in the window, so a return value of *n* indicates the scroll box lies adjacent to text line *n* in the window (numbered from zero).

```
scroll_by_wheel(int clicks, int per_click)
```

When you use a wheeled mouse like the Microsoft IntelliMouse, Epsilon calls the `scroll_by_wheel()` subroutine whenever you roll its wheel. (See the next section for information on what happens when you click the wheel, not roll it.) Epsilon provides the number of clicks of the wheel since the last time this function was called (which may be positive or negative) and the control panel setting that indicates the number of lines Epsilon should scroll on each click.

After calling this subroutine, Epsilon can then optionally generate a `WIN_WHEEL_KEY` key event. See page 565.

Mouse Panning

```
int mouse_panning;
int mouse_panning_rate(int percent, int slow, int fast)
```


The `mouse_panning` variable and the `mouse_panning_rate()` primitive work together to support panning and auto-scroll with the Microsoft IntelliMouse (or any other three button mouse). The EEL subroutine that receives clicks of the third mouse button sets `mouse_panning` nonzero to tell Epsilon to begin panning and record the initial position of the mouse.

Then the subroutine can regularly call `mouse_panning_rate()` to determine how quickly, and in what direction, to scroll. The parameter `percent` specifies the percentage of the screen the mouse has to travel to reach maximum speed (usually 40%). The parameter `slow` specifies the minimum speed in milliseconds per screen line (usually 2000 ms/line). The parameter `fast` specifies the maximum speed in milliseconds per screen line (usually 1 ms/line).

The `mouse_panning_rate()` primitive uses these figures, plus the current position of the mouse, to return the scroll rate in milliseconds per screen line. It returns a positive number if Epsilon should scroll down, a negative number to scroll up, or zero if Epsilon should not scroll.

See the previous section for information on what happens when you roll the wheel on a wheeled mouse instead of clicking it.

10.6.3 Window Events

When an EEL function calls `getkey()` to retrieve the next key, it sometimes receives a key code that doesn't correspond to any actual key, but represents some other kind of input event. Mouse keys (see page 560) are one example of this. This section describes the other key codes Epsilon uses for input events. These keys only occur in the Windows version.

The `WIN_MENU_SELECT` key indicates that the user selected an item from a menu or the tool bar. Epsilon sets the variable `menu_command` to the name of the selected command whenever it returns this key.

The `WIN_DRAG_DROP` key indicates that the user has just dropped a file on one of Epsilon's windows, or that Epsilon has received a DDE message from another program. See the description of the `drag_drop_result()` primitive on page 523.

The `WIN_EXIT` key indicates that the user has tried to close Epsilon, by clicking on the close box, for example.

The `WIN_HELP_REQUEST` key indicates that the user has just pushed a button in Epsilon's help file to set a particular variable or run a command. Epsilon fills the `menu_command` variable with the message from the help system.

The `GETFOCUS` and `LOSEFOCUS` keys indicate that a particular screen has gained or lost the focus. These set `mouse_screen` just like mouse keys. (See page 560.)

The `WIN_RESIZE` key indicates that Epsilon has resized a screen. Sometimes Epsilon will resize the screen without returning this key.

The `WIN_VERT_SCROLL` key indicates that Epsilon has scrolled a window. Epsilon doesn't normally return keys for these events. Instead, Epsilon calls the EEL subroutine `scrollbar_handler()` from within the `wait_for_key()` function, passing it information on which scroll bar was clicked, which part of the scroll bar was selected, and so forth.

Epsilon only recognizes user attempts to scroll by clicking on the scroll bar, or to resize the window, when it waits for a key in a recursive edit level. When an EEL command requests a key, Epsilon normally ignores attempts to scroll, and postpones acting on resize attempts.

An EEL command can set the `permit_window_keys` variable to allow these things to happen immediately, and possibly redraw the screen. Bits in the variable control these activities: set the `PERMIT_SCROLL_KEY` bit to permit immediate scrolling, and set `PERMIT_RESIZE_KEY` to permit resizing. Setting `PERMIT_SCROLL_KEY` also makes Epsilon return the `WIN_VERT_SCROLL` key shortly after scrolling.

Setting the `PERMIT_WHEEL_KEY` bit tells Epsilon to generate a `WIN_WHEEL_KEY` key event after scrolling due to a wheel roll on a Microsoft IntelliMouse.

The `WIN_BUTTON` key indicates that the user has clicked on a button in a dialog box, or selected the button via the keyboard. By default, Epsilon translates many buttons to standard keys like `Ctrl-m`. An EEL program can set the variable `return_raw_buttons` to disable this translation and instead receive `WIN_BUTTON` keys for each button pressed. For other buttons, and for check boxes and certain other dialog events, Epsilon always enqueues a `WIN_BUTTON` key. For each of these input events, it sets the `key_is_button` variable to a distinct value.

Epsilon supports up to three buttons on a mouse directly, with distinct key codes like `MOUSE_CENTER_DN` (plus wheel events). If Epsilon for X11 receives messages about additional buttons, it returns them using the `WIN_BUTTON` key code, with `key_is_button` set to a unique value for that button.

10.6.4 Completion Subroutines

There are several EEL subroutines defined in `complete.e` that get a line of input from the user, allowing normal editing. Most of them offer some sort of completion as well. They also provide a command history.

Each function takes two or three arguments. The first argument is an array of characters in which to store the result. The second argument is a prompt string to print in the echo area. The third argument, if there is one, is the default string. Depending on the setting of the `insert-default-response` variable, Epsilon may insert this string after the prompt, highlighted, or it may be available by pressing `Ctrl-r` or `Ctrl-s`.

Some functions will substitute the default string if you press `<Enter>` without typing any response. These functions display the default to you inside square brackets `[]` (whenever they don't actually pre-type the default after the prompt). The prompt that you must provide to these functions shouldn't include the square brackets, or the colon and space that typically ends an Epsilon prompt. The function will add these on before it displays the prompt. If there should be no default, use the empty string `""`.

```
get_file(char *res, char *pr, char *def)
get_file_dir(char *res, char *pr)
```

The `get_file()` and `get_file_dir()` subroutines provide file name completion. When the `get_file()` subroutine constructs its prompt, it begins with the prompt string `pr`, then appends a colon `:` and a space. (If `insert-default-response` is zero, it also includes the default value in the prompt, inside square brackets.) If the user presses `<Enter>` without typing any response, `get_file()` copies the default `def` to the response string `res`.

The `get_file_dir()` subroutine provides the directory part of the current file name, inserted as part of a default response or available via `Ctrl-s` or `Ctrl-r` (see the description of the `prompt-with-buffer-directory` variable), but it doesn't display that as part of the prompt. It uses the prompt `pr` as is. It doesn't substitute any default if the user enters no file name. Both `get_file()` and `get_file_dir()` call `absolute()` on the name of the file before returning (see page 510).

```
get_buf(char *res, char *pr, char *def)
```

The `get_buf()` subroutine completes on the name of a buffer. To construct its prompt, the subroutine begins with the prompt string `pr`, then adds the default `def` inside square brackets `[]`, and then appends a colon `:` and a space.

```
get_any(char *res, char *pr, char *def)
get_cmd(char *res, char *pr, char *def)
```

```

get_macname(char *res, char *pr, char *def)
get_func(char *res, char *pr, char *def)
get_var(char *res, char *pr, char *def, int flags)

```

Epsilon locates commands, subroutines, and variables by looking them up in its *name table*. See page 543 for details. The subroutines that complete on commands, variables and so forth all look in the same table, but restrict their attention to particular types of name table entries. For example, the `get_macname()` subroutine ignores all name table entries except those for keyboard macros. In the following table, ● indicates that the subroutine allows entries of that type.

	Command	Subr.	Kbd. Macro	Key Table	Variable
<code>get_any()</code>	●	●	●	●	●
<code>get_cmd()</code>	●		●		
<code>get_func()</code>	●	●			
<code>get_macname()</code>			●		
<code>get_var()</code>					●

These subroutines all substitute the default string if you just press `<Enter>` without entering anything. They also display the default inside square brackets `[]` after the prompt you provide (if `insert-default-response` is zero), and then append a colon `:` and a space.

The `get_var()` subroutine takes an additional, fourth parameter. It contains a set of flags to pass to the `comp_read()` subroutine, as listed below.

```
int get_command_index(char *pr)
```

The `get_command_index()` subroutine defined in `control.e` calls the `get_cmd()` subroutine to ask the user for the name of a command. It then checks to see if the command exists, and reports an error if it doesn't. (When checking, it allows subroutines and macros as well as actual commands.) If the function name checks out, `get_command_index()` returns its name table index.

Completion Internals

```

/* bits for finder func */
#define STARTMATCH      1
#define EXACTONLY       2
#define LISTMATCH       4
#define FM_NO_DIRS      (0x10)
#define FM_ONLY_DIRS    (0x20)
char *b_match(char *partial, int flags)
/* sample finder */

comp_read(char *response, char *prmt,
          char *(*finder)(), int flags, char *def)

/* bits for comp_read() */
#define CAUTIOUS        (0x100)
#define COMP_FOLD       (0x200)
#define MUST_MATCH      (0x400)
#define NONE_OK         (0x800)
#define POP_UP_PROMPT   (0x1000)

```

```

#define COMP_FILE      (0x2000 | CAUTIOUS)
#define PASSWORD_PROMPT (0x4000)
#define SPACE_VALID    (0x8000)

prompt_comp_read(char *response, char *prmt,
                 char *(*finder)(), int flags,
                 char *def)
zeroed char completion_column_marker;

```

It's easy to add new subroutines that can complete on other things. First, you must write a “finder” function that returns each of the possible matches, one at a time, for something the user has typed. For example, the `get_buf()` subroutine uses the finder function `b_match()`.

A finder function takes a parameter `partial` which contains what the user's typed so far, and a set of flags. If the `STARTMATCH` flag is on, the function must return the first match of `partial`. If `STARTMATCH` is off, it should return the next match. The function should return 0 when there are no more matches. The `LISTMATCH` flag is on when Epsilon is preparing a list of choices because the user has pressed '?'. This is so that a finder function can format the results differently in that case. If the `EXACTONLY` flag is on, the finder function should return only exact matches for `partial`. If the finder function is matching file names, you may also provide the `FM_NO_DIRS` flag, to exclude directory names, or `FM_ONLY_DIRS` to retrieve only directory names.

Next, write a subroutine like the various `get_` routines described above, all of which are defined in `complete.e`. It should take a prompt string, possibly a default string, and a character pointer in which to put the user's response. It passes these to the `comp_read()` subroutine, along with the name of your finder function (as a function pointer).

The `comp_read()` subroutine also takes a `flags` parameter. If the `CAUTIOUS` flag is zero, `comp_read()` assumes that all matches for a certain string will begin with that string, and that if there is only one match for a certain string, adding characters to that string won't generate any more matches. These assumptions are true for most things Epsilon completes on, but they're not true for files. (For example, if the only match for `x` is `xyz`, but `xyz` is a directory with many files, the second assumption would be false. The first assumption is false when Epsilon completes on wildcard patterns like `*.c`, since none of the matches will start with the `*` character.) If you provide the `CAUTIOUS` flag when you call `comp_read()`, Epsilon doesn't make those assumptions, and completion is somewhat slower.

Actually, when completing on files, provide the `COMP_FILE` macro instead of just `CAUTIOUS`; this includes `CAUTIOUS` but also makes Epsilon use some special rules necessary for completing on file names.

If you provide the `COMP_FOLD` flag to `comp_read()`, it will do case-folding when comparing possible completions.

The `MUST_MATCH` flag tells `comp_read()` that if the user types a response that the finder function doesn't recognize, it's probably a mistake. The `comp_read()` subroutine will then offer a list of possible responses, even though the user may not have pressed a key that ordinarily triggers completion. The `comp_read()` subroutine might still return with an unrecognized response, though. This flag is simply advice to `comp_read()`. The `NONE_OK` flag is used only with `MUST_MATCH`. It tells `comp_read()` that an empty response (just typing <Enter>) is ok.

Under Epsilon for Windows, the `POP_UP_PROMPT` flag tells `comp_read()` to immediately pop up a one-line dialog box when prompting. Right now, this flag may only be used when no completion is involved, and `comp_read()` is simply prompting for a line of text.

The `PASSWORD_PROMPT` flag tells `comp_read()` to display each character of the response as a `*` character. When the Internet functions prompt for a password they use this flag.

The `SPACE_VALID` flag tells `comp_read()` that a `<Space>` character is valid in the response. Since `<Space>` is also a completion character, `comp_read()` tries to guess whether to add a `<Space>` or complete, by examining possible matches.

A finder function receives any of the above flags that were passed to `comp_read()`, so it can alter its behavior if it wants.

The `comp_read()` subroutine uses the prompt you supply as-is. Usually, the prompt should end with a colon and a space, like "Find file: ". By contrast, the `prompt_comp_read()` subroutine adds to the supplied prompt by showing the default value inside square brackets, when `insert-default-response` is zero. The prompt string you supply to it should not end with a colon and space, since Epsilon will add these. If you provide a prompt such as "Buffer name" and a default value of "main", Epsilon will display "Buffer name [main]: ". If the default value you provide is empty or too long, Epsilon will instead display "Buffer name: ", omitting the default. Whether or not Epsilon displays the default, if the user doesn't enter any text at the prompt the `prompt_comp_read()` subroutine substitutes the default value by copying `def` to `response`.

Sometimes it's convenient if a finder function returns matches with additional data after them which shouldn't be included in the returned response. You can set the variable `completion_column_marker` to a character that marks the start of such extra data. If it's nonzero, and a response from the finder function contains it, Epsilon strips that character and any following text from the response before returning it, if the user selects that choice from the list of completions.

```
char>(*list_finder)();
list_matches(char *s, char>(*finder)(), int flags, int mbuf)
int>(*completion_lister)();
char resize_menu_list;
```

The `comp_read()` subroutine looks at several variables whenever it needs to display a list of possible completions (such as when the user types '?'). You can change the way Epsilon displays the list by setting these variables. Typically, you would use the `save_var` statement to temporarily set one of these while your completion routine runs.

By default, Epsilon calls the `list_matches()` subroutine to prepare its buffer of possible matches. The function takes the string to complete on, the finder function to use, flags as described above, and a buffer number. It calls the finder function repeatedly (passing it the `LISTMATCH` flag as well as any others passed to `list_matches()`) and puts the resulting matches into the indicated buffer, after sorting the matches. If the `completion_lister` function pointer is non-null, Epsilon calls that function instead of `list_matches()`, passing it the same parameters. If, for example, you have to sort the matches in a special order, you can set this variable.

If you simply want a different list of matches when Epsilon lists them, as opposed to when Epsilon completes on them, you can set the `list_finder` function pointer to point to a different finder function. The `list_matches()` subroutine always uses this variable if non-null, instead of the finder function it receives as a parameter.

An EEL completion function can temporarily set the `resize_menu_list` variable nonzero to indicate that if the user tries to list possible completion choices, the window displaying the choices should be widened if necessary to fit the widest choice. This variable has no effect on Epsilon windows within GUI dialogs.

```
int complete(char *response, char(*finder)(), int flags)
```

To actually do completion, `comp_read()` calls the `complete()` subroutine. It takes a finder function pointer, flags like `CAUTIOUS` and `COMP_FOLD` described above, and a string to complete on. It tries to extend the string with additional characters from the matches, modifying it in place.

The `complete()` subroutine generally returns the number of possible matches for the string. However, it may be able to determine that no more completion is possible before reaching the last match. For example, if the subroutine tries to complete on the file name “foo”, and encounters files named “foobar”, “foobaz”, “foo3”, “foo4” and so forth, it can determine on the third file that no completion is possible. In this case, it returns 3, even though there may be additional matches. It can only “give up early” in this way when it has encountered two or more matches. So when the subroutine returns a value of two or greater, there may be additional matches not included in its count.

```
build_prompt(char *full, char *pr, char *def, int omit, int rel)
```

The `build_prompt()` subroutine helps construct the text of a prompt. It copies the prompt *pr* to *full*, appending the default value *def* to it (inside brackets).

If the combination would be too wide for the screen, the subroutine abbreviates the default value. If even an abbreviated value would be too wide, or if *omit* is nonzero, it omits the default from the prompt entirely. If *rel* is nonzero, it assumes *def* is an absolute pathname, and uses its relative form.

```
find_buffer_prefix(int buf, char *prefix)
```

The `find_buffer_prefix()` subroutine looks through all the lines in the buffer *buf* to see if they all start with the same string of characters. It puts any such common prefix shared by all the lines in *prefix*. For instance, if the buffer contains three lines “waters”, “watering” and “waterfall”, it would put the string “water” in *dest*.

```
char *general_matcher(char *s, int flags)
```

Epsilon provides a general-purpose finder function called `general_matcher()`. An EEL function can perform completion on some arbitrary list of words by putting the list of words in a buffer named `_MATCH_BUF` (a macro defined in `eel.h`) and then providing `general_matcher()` as a finder function to a subroutine like `comp_read()`. Call `comp_read()` with the `COMP_FOLD` flag if you want `general_matcher()` to ignore case when comparing.

```
char _doing_input;
keytable comp_tab, menu_tab, view_tab;
```

An EEL function can tell if Epsilon is currently prompting for a line of input (or performing certain related tasks) by examining the `_doing_input` variable. It’s zero normally. While Epsilon is inside the `search_read()` subroutine or related ones, getting a search string from the user, it’s set to `DI_SEARCH`. While Epsilon is prompting for a line of some other type of input, it’s set to `DI_LINEINPUT`.

While the `view_buf()` subroutine (or a related one) is displaying a pop-up window, `_doing_input` is set to either `DI_VIEW` or `DI_VIEWLAST`, according to whether the `last` parameter passed to `view_buf()` was zero or not.

When Epsilon prompts for input, it uses certain specialized key tables in the buffer that accepts the input. It uses the `comp_tab` key table for the one-line buffer where the user types a response, in all the subroutines that accept a line of input (except in `search_read()` and related subroutines that prompt for search strings).

When Epsilon displays a list of possible matches, or previous responses, or similar things, while getting a line of input, it uses the `menu_tab` key table in the buffer displaying the list.

Finally, subroutines like `view_buf()` normally use the `view_tab` key table to show a buffer in a pop-up window.

Listing Commands, Buffers, or Files

```
int name_match(char *prefix, int start)
```

Several primitives help to perform completion. The `name_match()` primitive takes a command prefix such as “nex” and a number. It finds the next command or other name table entry that begins with the supplied prefix, returning its name table index. If its numeric argument is nonzero, it starts at the beginning of the name table. Otherwise it continues from the name table index returned on the previous call. It returns zero when there are no more matching names. When comparing names, case doesn’t count and ‘-’ is the same as ‘_’.

```
char *buf_match(char *pattern, int flags)
char *do_file_match(char *pattern, int flags)
char *low_file_match(char *pattern, int flags)
#define STARTMATCH      1
#define EXACTONLY       2
#define FM_NO_DIRS      (0x10)
#define FM_ONLY_DIRS    (0x20)
#define FM_FOLD         (0x200)
char *file_match(char *pattern, int flags)
```

The `buf_match()` and `file_match()` primitives may be used to find buffer or file names matching a pattern, similar to `name_match()`. Instead of returning a command index, they return the actual matching buffer or file names, respectively, and return a null pointer when there are no more matches. The values returned by `file_match()` are only valid until the next call to this function. Copy the name if you want to preserve it.

The `buf_match()` primitive returns one of a series of buffer names that match a pattern. The pattern is of the sort that `fpatternmatch()` accepts: * matches any number of characters, ? matches a single character, [a-z] represents a character class, and | separates alternatives. The `STARTMATCH` flag tells it to examine the pattern and return the first match; omitting the flag makes it return the next match of the current pattern. The `EXACTONLY` flag tells it to return only exact matches of the pattern; otherwise it returns buffer names that start with a match of the pattern (as if it ended in *). The `FM_FOLD` flag tells it to ignore case when comparing buffer names against the pattern; by default buffer names are case-sensitive (but see the `preserve-filename-case` variable).

The `file_match()` primitive returns one of a series of file names that match a pattern. You can use this primitive to expand file name patterns such as `a*.c`. See page 138 for details on Epsilon’s syntax for file patterns. The `STARTMATCH` flag tells it to examine the pattern and return the first match; omitting the flag makes it return the next match of the current pattern. The `EXACTONLY` flag tells it to return only exact matches of the pattern; otherwise it returns file names that start with a match of the pattern. Use the `FM_NO_DIRS` flags if you want to skip over directories when looking for files that match, or `FM_ONLY_DIRS` to retrieve only directory names.

Instead of directly calling the `file_match()` primitive, you may want to call one of the subroutines `do_file_match()` or `low_file_match()`. They both take the same arguments as `file_match()` and return the same values. But they also handle URL file patterns like `scp://host/*.c` (unless the `file-pattern-rules` variable includes the `FPAT_PERMIT_NO_URLS` bit). Additionally, `do_file_match()` removes any " double-quote characters from the pattern before using it.

```
short abort_file_matching = 0;
#define ABORT_IGNORE 0 /* ignore abort key & continue */
#define ABORT_JUMP -1 /* jump via check_abort() */
#define ABORT_ERROR -2 /* return ABORT_ERROR as error code */
```

By default, the `file_match()` and `do_dired()` primitives ignore the abort key. (See page 509 for information on `do_dired()`.) To permit aborting a long file match, set the primitive variable `abort_file_matching` using `save_var` to tell Epsilon what to do when the user presses the abort key. If you set `abort_file_matching` to `ABORT_ERROR` and the user presses the abort key, this function will return a failure code and set `errno` to `EREADABORT`. Set the variable to `ABORT_JUMP` if you want Epsilon to abort your function by calling the `check_abort()` primitive. (See page 534.) By default, the variable is zero, and Epsilon ignores the abort key until the primitive finishes.

10.6.5 Other Input Functions

```
get_strdef(char *res, char *pr, char *def)
get_strnone(char *res, char *pr, char *def)
get_string(char *res, char *pr)
get_str_auto_def(char *res, char *pr)
get_strpopup(char *res, char *title,
             char *def, char *help)
```

The subroutines `get_string()`, `get_strdef()`, and the rest each get a string from the user, and perform no completion. They each display the prompt, and accept a line of input with editing.

The `get_strdef()` routine additionally displays the default string (indicated by `def`) and allows the user to select the default by typing just the `<Enter>` key. The user can also pull in the default with `Ctrl-s`, and then edit the string if desired. While the other two functions use their prompt arguments as-is, `get_strdef()` constructs the actual prompt by adding a colon and space. If `insert-default-response` is zero, they also include the default value in the prompt, inside square brackets.

The `get_strnone()` subroutine works like `get_strdef()`, except that the default string is not displayed in the prompt (even when `insert-default-response` is zero), and Epsilon won't replace an empty response with the default string. Use this instead of `get_strdef()` if an empty response is valid.

The `get_str_auto_def()` subroutine is like `get_strdef()`, except it automatically provides the last response to the current prompt as a default.

The `get_strpopup()` subroutine is a variation of `get_strnone()` that is only available under Epsilon for Windows. It displays a simple dialog. The parameter `title` provides the dialog's title, and `def` provides the initial contents of the response area, which is returned in `res`. If the user presses the Help button, Epsilon will look up help for the specified command or variable name or other topic name in its help file.

```
int get_number(char *pr)
int numtoi(char *str)
int numtoi_with_base(char *str, int base)
int strttoi(char *str, int base)
int exptoi(char *str)
int evaluate_numeric_expression(char *expr)
char got_bad_number;
```

The `get_number()` subroutine is handy when a command needs a number. It prompts for the number using `get_string()`, but uses the prefix argument instead if one is provided. It returns the number obtained, and also takes care of resetting `iter` if necessary. It also understands numbers such as `0x10` in EEL's hexadecimal (base 16) format, binary and octal numbers, and character codes like `'a'`.

The `numtoi()` subroutine converts from a string to a number. It skips over any spaces at the beginning of its string parameter, determines the base (by seeing if the string starts with `"0x"` or similar), and then calls

`strtoi()` to perform the actual conversion. The subroutine `strtoi()` takes a string and a base, and returns the value of the string assuming it is a number in that base. It handles bases from 2 to 16, and negative numbers too. It stops when it finds a character that is not a legal digit in the requested base. The `numtoi_with_base()` subroutine is similar, but where `numtoi()` assumes base 10 numbers when the user doesn't specify a base, with `numtoi_with_base()` you can specify a default base.

The `evaluate_numeric_expression()` subroutine evaluates an arithmetic expression that may contain operators like `+`, `-`, `*`, or `/`, returning its numeric value. It runs the EEL compiler to do this. The `exptoi()` subroutine does the same. However, it examines the expression first. If it contains no arithmetic operators, it calls `numtoi()` instead, which is significantly faster.

The `numtoi()` and `exptoi()` subroutines both also recognize a key name inside angle brackets, such as `<Tab>`, and return the key's numeric value. (The `evaluate_numeric_expression()` subroutine doesn't, only permitting proper EEL expressions.) The `exptoi()` subroutine is usually the best choice for evaluating numeric user input. The `get_number()` subroutine mentioned above calls it.

All these subroutines set the variable `got_bad_number` to a nonzero value if the string they receive doesn't indicate a valid number. They return the value zero in this case. If the string does represent a number, they set `got_bad_number` to zero.

```
int get_choice(int list, char *resp, char *title,
               char *msg, char *b1, char *b2,
               char *b3)
int get_key_choice(int list, char *resp, char *title,
                  char *msg, char *b1, char *b2, char *b3)
int select_menu_item(int resbuf, int menuwin,
                    int owin, int dir)
```

The `get_choice()` subroutine provides a way to ask the user to select one of a list of choices. The choices must appear in the buffer `list`, one to a line. The subroutine displays a pop-up window with the indicated title and shows the specified message.

Epsilon for Windows instead displays a dialog with the indicated title, and doesn't use the message. It uses the specified button labels (see the description of the `button_dialog()` primitive on page 575 for details). The `get_choice()` subroutine puts the user's choice in `resp` and returns 1. If the user cancels, the subroutine returns 0.

If `resp` is initially nonempty, `get_choice()` will position point on the first line starting with that text. If `resp` is initially "", the subroutine won't change point in `list`.

The `get_key_choice()` subroutine is a variation on `get_choice()`, designed to let the user select one of a series of options. Each of the lines in the buffer `list` should begin with a unique character. When the user presses that character, Epsilon moves to that entry in the list. This subroutine's behavior is otherwise identical to `get_choice()`.

The `get_choice()` subroutine uses the `select_menu_item()` subroutine to handle user interaction. It takes the window handle `menuwin` of a window containing a list of choices and returns when the user has selected one. The parameter `owin` should be the handle of the window that was current before displaying `menuwin`. If `resbuf` is nonzero, Epsilon will copy the selected line to the specified buffer.

The parameter `dir` tells Epsilon how to behave when the user presses self-inserting keys like 'a'. If `dir` is zero, the subroutine interprets N and P to move forward and back, and Q to quit. Other normal keys are ignored. If `dir` is 1 or -1, and `search-in-menu` is nonzero, normal keys are added to the result, and Epsilon searches for the first (if 1) or last (if -1) item that matches.

10.6.6 Dialogs

Standard Dialogs

```
short common_file_dlg(char *fname, char *title,
                    int *flags, int save,
                    ?char *filt_str, ?char *cust_filter,
                    ?int *filt_index)
short use_common_file_dlg(char *fname, char *title,
                        int *flags, int save)
int use_common_file_dialog()
```

Under Windows, the `common_file_dlg()` primitive displays the Common Open/Save File Dialog. The `fname` parameter should be initialized to the desired default file name; on return it will hold the file name the user selected. The `title` parameter specifies the title of the dialog window. Epsilon passes the `flags` parameter to Windows; definitions for useful flag values appear in `codes.h`. Windows modifies some of the flags before it returns from the dialog. If the parameter `save` is nonzero, Epsilon displays the Save dialog, if zero it uses the Open dialog. This primitive uses the `common-open-curdir` variable to hold the directory that this dialog should display.

The filter parameters let you specify the file types the user can select; these are all passed directly to Windows. Epsilon normally invokes `common_file_dlg()` through the `use_common_file_dlg()` subroutine, which uses the filter definitions from the `file.filter.txt` to construct the `filt_str` parameter. You can edit that file to add new filters.

The parameter `filt_str` has the following format. It consists of pairs of null-terminated strings. The first string says what to display in the dialog, while the second is a Windows-style list of file patterns, separated by semicolons. For example, the first string might be "Fortran files" and the second string might be "*.for;*.f77". A final null character follows the last pair.

The `use_common_file_dialog()` subroutine examines the `want-common-file-dialog` variable and other settings and tells whether a command should use the common file dialog in place of Epsilon's traditional file dialog.

```
find_dialog(int show)
find_dialog_say(char *text)
```

Under Windows, the `find_dialog()` primitive displays a find/replace dialog, when its parameter `show` is nonzero. When its parameter `show` is zero, it hides the dialog. While a find/replace dialog is on the screen, the `getkey()` function returns certain predefined keys to indicate dialog events such as clicking a button or modifying the search string. The `_find()` subroutine defined in `search.e` interprets these key codes to control the dialog. The global variable `find_data` lets that subroutine control the contents of the dialog.

When a find/replace dialog is on the screen, an EEL program can display an error message in it using the `find_dialog_say()` primitive. This also adds an alert symbol to the dialog. To clear the message and remove the alert symbol, pass a parameter of "".

```
short window_lines_visible(int w)
```

The `window_lines_visible()` primitive returns the number of lines of a given window that are visible above a find/replace dialog. If the given window contains twelve lines, but a find/replace dialog covers the bottom three, this function would return nine. If Epsilon isn't displaying a find/replace dialog, the function returns the number of lines in the given window.

```
int comm_dlg_color(int oldcolor, char *title)
```

In Epsilon for Windows, the `comm_dlg_color()` primitive lets the user select a color using the Windows common color dialog. The `oldcolor` parameter specifies the default color, and `title` specifies the dialog's title. The primitive returns the selected color, or -1 if the user canceled.

```
about_box()
```

The `about_box()` primitive displays Epsilon's "About" box under Windows. In other versions of Epsilon, it inserts similar information into the current buffer. The `about-epsilon` command uses this primitive.

Button Dialogs

```
short button_dialog(char *title, char *question,
                   char *yes, char *no, char *cancel,
                   int def_button)
```

The `button_dialog()` primitive displays a dialog having one to three buttons. By convention, these buttons have meanings of "Yes", "No", and "Cancel", but the labels may have any text. Set the `cancel` parameter to "" to use a dialog with two buttons. Set both `cancel` and `no` to "" if you want a dialog with one button. Put & before a character in a button label to make it an access key; it will be underlined, and pressing the key will act like clicking that button. Use && for a literal & character. The parameter `title` specifies the title of the dialog. The parameter `question` holds the text to display in the dialog next to the buttons.

Set `def_button` to 1, 2, or 3 to make the default button be the first, second or third. Any other value for `def_button` is the same as 1. Canceling or closing the dialog is equivalent to pressing the last defined button.

The primitive returns 1, 2, or 3 to indicate which button was pressed. It sets the `key_is_button` variable to its return value if the user actually clicked a button, or to zero if he pressed a key to end the dialog. It sets `key` and `full_key` to indicate the pressed key (uppercased), or to 'Y', 'N', or 'C', respectively, if the user clicked one of the buttons.

If the user clicked the dialog's close box, the primitive returns 3, setting `key_is_button` to 4, and `key` and `full_key` to the abort key. If the user pressed the abort key, the primitive returns the code for the last button, setting `key_is_button` to 0, and `key` and `full_key` to the pressed key.

This primitive only works in the Windows version of Epsilon; read on for similar functions that work everywhere.

```
int ask_yn(char *title, char *question, char *yes_button,
           char *no_button, int def_button)
ask_3way(char *title, char *question, char *prompt,
         char *a1, char *a2, char *a3, int def)
```

The `ask_yn()` subroutine defined in `basic.e` asks a Yes/No question. Under Windows, it uses a dialog. The parameters specify the title of the dialog, the text of the question displayed in it, and the text on its two buttons (typically "Yes" and "No", but sometimes "Save" and "Cancel" or the like). Put & before a character in a button label to make it an access key; it will be underlined, and pressing the key will act like clicking that button. Use && for a literal & character.

Set `def_button` to 0 for no default, 1 to make the first choice “Yes” the default, or 2 to make the second choice “No” the default. (Under non-Windows versions, no default means that just hitting `<Enter>` won’t return from this function; you must choose an option. Under Windows, no default is the same as a default of “Yes”.) The function returns 1 if the user selected the first option “Yes” or 0 if the user selected the section option “No”. Non-Windows versions of Epsilon only use the `question` and `def_button` parameters. They modify the prompt to indicate the default, if any.

The `ask_3way()` subroutine is similar, but asks a question with three possible responses. Epsilon always displays the prompt; GUI versions also display a dialog with the specified title and question. The first letter of each answer is its shortcut key; put an ampersand “&” before another letter in an answer to use that as the shortcut key instead. The value `def` may be 1, 2, or 3 to make that answer the default, or 0 for none. The subroutine returns 1, 2, 3 as its response, or 0 if the user aborted.

```
int get_key_response(char *pr, char *valid, int def, char *helpcmd)
```

The `get_key_response()` subroutine waits for the user to type a valid key in response to a prompt `pr`. The parameter `valid` lists the acceptable characters, such as “YN” for a yes/no question. (But see the `ask_yn()` subroutine, more suitable for yes/no questions.) The `def` parameter, if greater than zero, indicates which key should be the default if the user presses `<Enter>`. If the `helpcmd` parameter is non-null, Epsilon displays help on that topic if the user presses `?` or another help key. The subroutine returns the selected key.

Windowed Dialogs

```
display_dialog_box(char *dialogname, char *title,
                  int win1, int win2, int win3,
                  char *button1, char *button2, char *button3,
                  char *button4, ?char *tag)
```

The `display_dialog_box()` primitive creates a new dialog box in Epsilon for Windows containing one or more Epsilon windows. The `dialogname` must correspond to one of the dialogs in this list:

Dialog name	Windows	Dialog name	Windows
AskExitBox	2	GeneralBox	1
AskSaveBox	2	HelpSetup1	1
CaptionBox	2	OneLineBox	1
EditVarBox	2	PromptBox	2
FileDateBox	1	SetColorBox	3
FileDateBox2	1	UsageBox	1

Each dialog requires one to three handles to pop-up windows, created with `add_popup()` in the usual way. The primitive moves these windows to the new dialog box. If you use a dialog which requires only one or two window handles, provide zero for the remaining handles. The windows will be resized to fit the dialog, and each will be assigned a unique “screen handle”. Mouse clicks in that window will set the `mouse_screen` variable to the matching screen handle. You can use the `window_to_screen()` primitive to determine the screen number assigned to each window.

The parameters `button1`, `button2`, `button3`, and `button4` specify the text for the buttons. If you want fewer buttons, provide the value “” for any button except `button1` and it will not appear. The specified `title` appears at the top of the dialog box.

When you click on a button in a dialog, Epsilon normally returns a particular fixed keystroke: either `Ctrl-m`, or the abort key specified by the `abort_key` variable, or the help key specified by the `HELPKEY`

macro, for the first, second, and third buttons respectively. These correspond to typical button labels of “OK”, “Cancel”, and “Help”, so that most EEL programs don’t need to do anything special to receive input from buttons. If an EEL program needs to know whether a keypress came from an actual key, or a button, it can examine the value of the `key_is_button` variable. This variable is zero whenever the last key returned was an actual key, and nonzero when it was really a button. In the latter case, its value is 1 if the leftmost button was pressed, 2 if the next button was pressed, and so forth.

Sometimes an EEL program puts different labels on the buttons. It can be more convenient in this case to retrieve a button press as a distinct key. Set the `return_raw_buttons` variable to a nonzero value to retrieve all button presses as the key code `WIN_BUTTON`. The `key_is_button` variable will still be set as described above, so you can distinguish one button from another by examining its value.

Epsilon for Windows remembers and restores the sizes of most of these dialogs. Some dialogs may be used in multiple contexts. To have Epsilon remember a different size for each context, pass a unique tag parameter. If the optional tag is missing or `NULL`, Epsilon uses a default context for remembering the dialog’s size. You can use `save_var` to temporarily assign a value to the `use_alternate_dialog_tag` variable, and the various subroutines below that call `display_dialog_box()` will use it as the dialog’s tag.

```
one_window_to_dialog(char *title, int win1,
                    char *button1, char *button2, char *button3)
prompt_box(char *title, int win1, int win2, char *tag)
two_scroll_box(char *title, int win1, int win2,
              char *button1, char *button2, char *button3)
void (*use_alternate_dialog)(int win1, int win2, int win3);
char *use_alternate_dialog_name;
```

The subroutines `one_window_to_dialog()`, `prompt_box()`, and `two_scroll_box()` each call `display_dialog_box()` with some of its parameters filled in for you. They display certain common kinds of dialogs. Call `one_window_to_dialog()` to display a dialog with a single text window and one to three buttons. To see an example, define a bookmark with Alt-/ and then type Alt-x list-bookmarks. Call `prompt_box()` to display a dialog with a one-line window, and below it a list-box style window. To see an example, type Ctrl-x Ctrl-f and then ‘?’. Call `two_scroll_box()` to display a dialog box with two multi-line windows.

These subroutines all call a subroutine named `do_display_dialog_box()`, which takes the same parameters as `display_dialog_box()`, but can be told to use an alternative function to display the dialog, not `display_dialog_box()`, by temporarily setting the function pointer `use_alternate_dialog` to a suitable function. Or they can be told to use a different dialog name by temporarily setting the `use_alternate_dialog_name` variable to its name.

```
next_dialog_item()
prev_dialog_item()
```

Within an Epsilon window that’s part of a dialog box, the `next_dialog_item()` and `prev_dialog_item()` primitives move the focus to a different window or button within the dialog box. Epsilon normally binds <Tab> and Shift-<Tab> to commands that use these primitives.

```
int dialog_checkboxes;
int disable_dialog_controls;
```

Some dialogs include check boxes. The `dialog_checkboxes` variable controls which check boxes are checked. Each check box corresponds to a bit in this variable. When EEL code sets or clears a bit, it’s

reflected in the state of that check box; similarly, when the user clicks the check box, its corresponding bit in this variable changes to match. This also generates a `WIN_BUTTON` event; see page 566.

The `disable_dialog_controls` variable lets EEL code disable check boxes and buttons in a dialog. Each check box or button corresponds to a bit in this variable. Setting the bit disables the corresponding control. Clearing the bit enables the control. For check boxes, it unchecks the check box, clearing the corresponding bit in `dialog_checkboxes`. Check boxes are assigned the low-order eight bits. Any buttons in the dialog are assigned the remaining bits.

```
set_window_caption(int win, char *title)
show_window_caption()
```

The `set_window_caption()` primitive sets the text in the title bar of the dialog box containing the window `win`. If the specified window is on Epsilon's main screen, it sets the main window title displayed above the menu bar. The `show_window_caption()` subroutine calls this to include the current file name in the caption of Epsilon's main window.

10.6.7 The Main Loop

While Epsilon runs, it repeatedly gets keys, executes the commands bound to them, and displays any changes to buffers that result. We call this process the *main loop*. Epsilon loops until you call the `leave_recursion()` primitive, as described on page 535. The steps in the main loop are as follows:

- Epsilon resets the `in_echo_area` variable. See page 481.
- Epsilon calls the `check_abort()` primitive to see if you pressed the abort key since the last time `check_abort()` was called. If so, an abort happens. See page 534.
- Epsilon sets the current buffer to be the buffer connected to the current window.
- Epsilon calls `maybe_refresh()`, so that all windows are brought up to date if the next key is not ready yet.
- Epsilon calls `undo_mainloop()`, to make sure undo information is kept for the current buffer, and to tell the undo system that future buffer changes will be part of the next command.
- Epsilon sets the `this_cmd` and `has_arg` variables to 0, and the `iter` variable to 1. See below.
- Epsilon calls the EEL subroutine `getkey()`. This subroutine in turn calls the `wait_for_key()` primitive to wait for the next key, mouse click, or other event.
- Epsilon executes the new key by calling the primitive `do_topkey()` as described on page 581.
- Epsilon sets the `prev_cmd` variable to the value in `this_cmd`.

```
user short this_cmd;
user short prev_cmd;
invisible_cmd()
```

Some commands behave differently depending on what command preceded them. For example, up-line behaves differently when the previous command was also up-line. To get this behavior, the command acts differently if `prev_cmd` is set to a certain value and sets `this_cmd` to that value itself. Epsilon copies the value in `this_cmd` to `prev_cmd` and then clears `this_cmd`, each time through the main loop.

Sometimes a command doesn't wish to be counted when determining the previous command. For example, when you move the mouse, Epsilon is actually running a command. But the up-line command of the previous example must behave the same, whether or not you happen to move the mouse between one up-line and the next. A command may call the `invisible_cmd()` primitive to make commands like up-line ignore it. (In fact, the primitive simply sets `this_cmd` equal to `prev_cmd`.)

```
user char has_arg;
user int iter;
```

Numeric arguments work using the `has_arg` and `iter` variables. The main loop resets `iter` to 1 and `has_arg` to 0. The argument command sets `iter` to the value of the argument, and sets `has_arg` to 1 so other commands can distinguish an argument of 1 from no argument. The `do_command()` primitive, described on page 581, will repeatedly execute a command while `iter`'s value is greater than one, subtracting one from `iter`'s value with each execution. If a command wants to handle arguments itself, it must set `iter` to one or less before returning, or the main loop will call it again.

```
user short cmd_len;
```

Any command may get more keys using the `wait_for_key()` primitive (usually by calling `getkey()`; see page 555). Epsilon counts the keys used so far by the current command and stores the value in the variable `cmd_len`. This counter is reset to zero each time Epsilon goes through the main loop. The counter doesn't count mouse keys or other events that appear as keys.

10.6.8 Bindings

Epsilon lets each buffer have a different set of key bindings appropriate to editing the type of text in that buffer. For instance, while in a buffer with EEL source code, a certain key could indent the current function. The same key might indent a paragraph in a buffer with text.

A key table stores a set of key bindings. A key table acts like a very large array, with one entry for each key on the keyboard. Each entry in the array contains an index into the name table. (See page 543.) If the value of a particular entry is negative or zero, it means the key is undefined according to that table.

Since the key codes that index a key table can be very large numbers, with big gaps between entries, Epsilon doesn't actually store key tables as arrays, but they act like arrays in EEL. (Internally, Epsilon stores a key table as a type of sparse array.) The `MAXKEYS` macro holds a number bigger than the biggest possible key code.

```
#define key_t      int
set_range(key_t *table, int i, int value, int cnt)
int get_range(key_t *table, int i, int value)
int find_next_entry(key_t *table, int value, int bound)
```

EEL code that wants to perform some operation on groups of keys can't simply use a loop like `for (i = 0; i < MAXKEYS; i++)` to do so; that would be too slow. Instead, there are primitives specifically designed for such purposes.

To set a range of key codes entries to a particular value, use the `set_range()` primitive. It sets the `cnt` entries starting at index `i` in the key table `table` to the specified value. The `key_t` type represents the type of a key table entry.

To find the length of a run of identical values in a key table array, use the `get_range()` primitive. It returns the number of entries, starting at index `i`, with the specified value. For instance, if `k` is a key table,

and `k[5]`, `k[6]`, and `k[7]` equal 123, but `k[8]` equals 456, then `get_range(k, 5, 123)` would return 3, and `get_range(k, 8, 123)` would return 0.

The `find_next_entry()` primitive searches for the index of the next member of the key table array `k` with the specified value. It starts looking at position `bound` in the array, skipping past index entries less than or equal to `bound`. It returns -1 if there are no more entries set to the specified value.

The EEL header file `oldkeys.h` provides sample implementations of the above primitives. If the number of possible keys were much smaller, you could use them instead of the above primitives. If you write EEL source code that must also work in older versions of Epsilon, you can include that file. Your EEL source code will use the definitions in `oldkeys.h` when you compile it in old versions of Epsilon; under Epsilon 12 or later, it will use the above primitives.

```
buffer short *mode_keys;
short *root_keys;
keytable reg_tab, c_tab;
```

Epsilon uses two key tables in its search for the binding of a key. First it looks in the key table referenced by the buffer-specific variable `mode_keys`. If the entry for the key is negative, Epsilon considers the command unbound and signals an error. If the entry for the key is 0, as it usually is, Epsilon uses the entry in the key table referenced by the variable `root_keys` instead. If the resulting entry is zero or negative, Epsilon considers the key unbound. If it finds an entry for the key that is a positive number, Epsilon considers that number the key's binding. The number is actually an index into the name table.

Most entries in a key table refer to commands, but an entry may also refer to a subroutine (if it takes no arguments), to a keyboard macro, or to another key table. For example, the entry for Ctrl-x in the default key table refers to a key table named `cx_tab`, which contains the Ctrl-x commands. The entry for the find-file command bound to Ctrl-x Ctrl-f appears in the `cx_tab` key table.

Normally in Epsilon the `root_keys` variable points to the `reg_tab` array. The `mode_keys` variable points to one of the many mode-specific tables, such as `c_tab` for C mode.

```
int new_table(char *name)
int make_anon_keytable()          /* control.e */
short *index_table(int index)
```

Key tables are usually defined with the `keytable` keyword as described on page 424. If a key table's name is not known when the routine is compiled, the `new_table()` primitive can be used. It makes a new key table with the given name. All entries in it are 0.

The `make_anon_keytable()` subroutine defined in `control.e` calls `new_table()`, first choosing an unused name for the table. The `index_table()` function takes a name table index and retrieves the key table it refers to.

```
fix_key_table(short *ftab, int fval, short *ttab, int tval)
copy_key_table(short *from, short *to)
set_list_keys(short *tab)
```

The `fix_key_table()` subroutine copies selected key table information from one key table to another. For each key in `ftab` bound to the function `fval`, the subroutine binds that key in `ttab` to the function `tval`. The `copy_key_table()` subroutine copies an entire key table. The `set_list_keys()` subroutine sets the 'n' and 'p' keys to move up or down by lines.


```
do_topkey()
run_topkey()
```

When Epsilon is ready to execute a key in its main loop, it calls the primitive `do_topkey()`. This primitive searches the key tables for the command bound to the current key, as described above. When it has found the name table index, it calls `do_command()`, below, to interpret the command.

The `run_topkey()` subroutine provides a wrapper around `do_topkey()` that resets `iter` and similar variables like the main loop does. An EEL subroutine that wants to retrieve keys itself and execute them as if the user typed them at command level can call this subroutine.

```
do_command(int index)
user short last_index;
```

The `do_command()` primitive executes the command or other item with the supplied name table index. If the index is invalid, then the `quick_abort()` primitive is called. Otherwise, the index is copied to the `last_index` variable, so the help system can find the name of the current command (among other uses).

If the name table index refers to a command or subroutine, Epsilon calls the function. When it returns, Epsilon checks the `iter` variable. If it is two or more, Epsilon proceeds to call the same function repeatedly, decrementing `iter` each time, so that it calls the function a total of `iter` times. See page 579.

```
key_t *table_keys;
int table_count;
table_prompt()                /* control.e */
```

If the entry in the name table that `do_command()` is to execute contains another table, Epsilon gets another key. First, Epsilon updates the primitive array `table_keys`. It contains the prefix keys entered so far in the current command, and `table_count` contains their number. Next, Epsilon calls the EEL subroutine `table_prompt()` if it exists to display a prompt for the new key. The version of this subroutine that's provided with Epsilon uses `mention()`, so the message may not appear immediately. Epsilon then calls the EEL subroutine `getkey()` to read a new key and clears the echo area of the prompt. Epsilon then interprets the key just as the `do_topkey()` primitive would, but using the new key table. If both `mode_keys` and `root_keys` provided a table as the entry for the first key, the values from each are used as the new mode and root key tables.

```
do_again()
```

The `do_again()` primitive reinterprets a key using the same pair of mode and root tables that were used previously. The value in the variable `key` may, of course, be different. Epsilon uses this primitive in commands such as `alt-prefix`.

Epsilon handles EEL subroutines without parameters in the name table in the same way as commands, as described above. If the entry is for a keyboard macro, the only other legal name table entry, Epsilon goes into a recursive edit level and begins processing the keys in the macro. It saves the macro internally so that future requests for a key will return characters from the macro, as described on page 555. It also saves the value of `iter`, so the macro will iterate properly. When the macro runs out of keys, Epsilon automatically exits the recursive edit level, and returns from the call to `do_again()`. (When `macro-runs-immediately` is nonzero, running a macro doesn't enter a recursive edit level, but returns immediately. Future key requests will still come from the macro until it's exhausted.)

```
short ignore_kbd_macro;
```

Epsilon provides a way for a keyboard macro to suspend itself and get input from the user, then continue. Set the `ignore_kbd_macro` variable nonzero to get keyboard input even when a macro is running. The `pause-macro` command uses this variable.

```
short *ask_key(char *pr, char *keyname, int flags)
int key_binding[30];      // ask_key() puts key info here
```

The `ask_key()` subroutine defined in `basic.e` duplicates the logic of the main loop in getting the sequence of keys that make up a command. However, it prompts for the sequence and doesn't run the command at the end. Commands like `bind-to-key` that ask for a key and accept a sequence of key table keys use it.

The `ask_key()` subroutine returns a pointer to the entry in the key table that was finally reached. The value pointed to is the name table index of the command the key sequence invokes.

This subroutine stores the key sequence in the `keyname` parameter in text form (as "Ctrl-x f", for example). It also copies the key sequence into the global variable `key_binding`. The key sequence is in macro format, so in the example of Ctrl-x f, `key_binding[1]` would hold CTRL('X'), `key_binding[2]` would hold 'f', and `key_binding[0]` would hold 3, the total number of entries in the array.

If you pass the 1 flag in `flags` and the user presses a key like <Backspace> with both a generic and a specific interpretation, Epsilon asks the user which one he wants. Without this flag, Epsilon assumes the specific key is meant.

If you pass the 2 flag and the user presses a key that normally shouldn't be rebound because it self-inserts (such as letter keys or <Enter>), the subroutine asks for confirmation.

```
full_getkey(char *pr, int code)      /* basic.e */

    /* for full_getkey() */
#define ALTIFY_KEY      1
#define CTRLIFY_KEY     2
```

The `full_getkey()` subroutine defined in `basic.e` gets a single key from the keyboard, but recognizes the prefix keys <Esc> and Ctrl-~. The `ask_key()` subroutine uses it, as well as the commands bound to the prefix keys above. It takes a prompt to display and a bit pattern (from `eel.h`) to make it act as if certain of the above keys had already been typed. For example, the `ctrl-prefix` command calls this subroutine with the value `CTRLIFY_KEY`. It leaves the key that results in the key primitive.

```
name_macro(char *name, key_t *keys)
```

Epsilon has no internal mechanism for capturing keyboard keys to build a macro (this is done in the `getkey()` subroutine defined in `control.e`), but once a macro has been built Epsilon can name it and make it accessible with the `name_macro()` function. It takes the name of the macro to create, and the sequence of keys making up the macro in an array of short ints. This array is in the same format that `get_keycode()` uses. That is, the first element of the array contains the number of valid elements in the array (including the first one). The actual keys in the macro follow. The `name_macro()` primitive makes a copy of the macro it is given, so the array can be reused once the macro has been defined.

```
key_t *get_macro(int index)
```

The `get_macro()` primitive can retrieve the keys in a defined keyboard macro. It takes the name table index of a macro, and returns a pointer to the array containing the macro, in the same format as `name_macro()`.

```
bind_universally(int k, int f)
```

The `bind_universally()` subroutine can be useful to bind a function to a key in all modes. When a key has a generic version, a mode that binds the generic version will override a global definition for the non-generic version of the key. For instance, the Tab key has a generic version, Ctrl-i. By default, each time you press Tab, Epsilon converts it to Ctrl-i and uses the current mode's binding for Ctrl-i, unless the current mode has its own definition for the Tab key too. If you want Ctrl-i to behave in a mode-specific manner, but force Tab to always run the same command regardless of the mode, you can call this function. It binds the key `k` to the function `f` in all key tables.

```
int list_bindings(int start, short *modetable,
                  short *roottable, int find)
```

The `list_bindings()` primitive quickly steps through a pair of key tables, looking for entries that have a certain name table index. It takes mode and root key tables, the name table index to find, and either -1 to start at the beginning of the key tables, or the value it returned on a previous call. It returns the index into the key table, or -1 if there are no more matching entries. For each position in the tables, Epsilon looks at the value in the mode key table, unless it is zero. In that case, it uses the root table.

In addition to the matches, `list_bindings()` also stops on each name table index corresponding to a key table, since these must normally be searched also. For example, the following file defines a command that counts the number of separate bindings of any command.

```
#include "eel.h"

command count_bindings()
{
    char cmd[80];

    get_cmd(cmd, "Count bindings of command", "");
    if (*cmd)
        say("The %s command has %d bindings", cmd,
            find_some(mode_keys,
                      root_keys, find_index(cmd)));
}

/* count bindings to index in table */
int find_some(modetable, roottable, index)
    short *modetable, *roottable;
{
    int i, total = 0, found;

    i = list_bindings(-1, modetable, roottable, index);
    while (i != -1) {

        found = (modetable[i]
                 ? modetable[i] : roottable[i]);
        if (found == index)
            total++;
        else
            total += find_some(index_table(found),
```

```

        index_table(found), index);

    i = list_bindings(i, modetable, roottable, index);
}
return total;
}

```

10.7 Defining Language Modes

There are several things to be done to define a new mode. Copy the file `samplemode.e`, included with Epsilon's EEL source code files, to get started.

Suppose you wish to define a mode called `reverse-mode` in which typing letters inserts them backwards, so typing “abc” produces “cba”, and yanking characters from a kill buffer inserts them in reverse order. First, define a key table for the mode with the `keytable` keyword, and put the special definitions for that mode in the table:

```

keytable rev_tab;

command reversed_normal_character()
{
    normal_character();
    point--;
}

when_loading()
{
    int i;

    for (i = 'a'; i <= 'z'; i++)
        rev_tab[toupper(i)] = rev_tab[i] = (short)
            reversed_normal_character;
}

command yank_reversed() on rev_tab[CTRL('Y')]
{
    ...
}

```

Now define a command whose name is that of the mode. It should set `mode_keys` to the new table and `major_mode` to the name of the mode, and then call the subroutine `make_mode()` to update the mode line:

```

command reverse_mode()
{
    mode_keys = rev_tab;          /* use these keys */
    major_mode = strkeep("Reverse");
    make_mode();
}

```

Using `strkeep()` for the mode name ensures that it remains valid even if the `reverse-mode` command is redefined later. Since some buffers may continue to point to it, it's important that the pointer remains valid. (Alternatively, you could define a character array variable with the mode name, and set `major_mode` to that.) The mode name in `major_mode`, with the addition of “-mode”, should be the name of a command that goes into that mode.

If you want Epsilon to go into that mode automatically when you find a file with the extension `.rev` (as it goes into C mode with `.c` files, for instance), define a function named `suffix_rev()` which calls `reverse_mode()`. The EEL subroutine `find_it()` defined in `files.e` automatically calls a function named `suffix_ext` (where *ext* is the file's extension) whenever you find a file, if a function with that name exists. It tries to call the `suffix_none()` function if the file has no suffix. If it can't find a function with the correct suffix, it will try to call the `suffix_default()` function instead.

```
suffix_rev()
{
    reverse_mode();
}
```

The source file `samplemode.e` defines a sample mode you can use as a template for your modes. Make a copy of the file, replace all references to “sample” with the name of your mode, and modify it as needed for your language's syntax.

Language modes may wish to define a compilation command. This tells the `compile-buffer` command on Alt-F3 how to compile the current buffer. For example, `compile_asm_cmd` is defined as `m1 "%r"`. (Note that " characters must be quoted with \ in strings.) Use one of the % sequences shown on page 122 in the command to indicate where the file name goes, typically %f or %r.

The mode can define coloring rules. See page 488 for details. Often, you can copy existing syntax coloring routines like those for `.asm` or `.html` files and modify them. They typically consist of a loop that searches for the next “interesting” construct (like a comment or keyword), followed by a `switch` statement that provides the coloring rule for each construct that could be found. Usually, finding an identifier calls a subroutine that does some additional processing (determining if the identifier is a keyword, for instance).

A language mode should set comment variables like `comment-start`. This tells the commenting commands (see page 104) how to search for and create legal comments in the language.

The comment commands look for comments using regular expression patterns contained in the buffer-specific variables `comment-pattern` (which should match the whole comment) and `comment-start` (which should match the sequence that begins a comment, like `/*`). When creating a comment, comment commands insert the contents of the buffer-specific variables `comment-begin` and `comment-end` around the new comment.

SHOWING MATCHING DELIMITERS.

Commands like `forward-level` that move forward and backward over matching delimiters will (by default) recognize `(`, `[`, and `{` delimiters. It won't know how to skip delimiters inside quoted strings, or similar language-specific features. A language mode can define a replacement delimiter movement function. See page 452 for details.

To let Epsilon automatically highlight matching delimiters in the language when the cursor appears on them, a language mode uses code like this to set the `auto-show-matching-characters` variable:

```
if (auto_show_asm_delimiters)
    auto_show_matching_characters = asm_auto_show_delim_chars;
```

where references to “asm” are of course replaced by the mode's name. The language mode should define the two variables referenced above:

```

user char auto_show_asm_delimiters = 1;
user char asm_auto_show_delim_chars[20] = "{[]}" ;

```

The list of delimiters should contain an even number of characters, with all left delimiters in the left half and right delimiters in the right half. (A delimiter that's legal on the left or right should appear in both halves; then the language must provide a `mode_move_level` definition that can determine the proper search direction itself. See page 452.)

Sometimes a mode may wish to highlight delimiters more complicated than single characters, such as BEGIN and END keywords. To do this, the mode should define a function such as `mymode_auto_show_delimiter()` and then set the buffer-specific function pointer variable `mode_auto_show_delimiter` to point to it in that buffer.

Epsilon will then call that function when idle to highlight delimiters. It should return 0 if no highlighting should be done, 1 to make Epsilon try to use the `auto_show_matching_characters` setting described above for simple highlighting, 2 to indicate mismatched delimiters, or 3 to indicate matched delimiters. In the latter two cases it should also display the highlighting, by setting two arrays to mark the appropriate buffer regions, as shown in the example. This sample only demonstrates how to control the highlighting; a typical mode would use smarter rules for finding the matching keywords (ignoring nested pairs, skipping over keywords in comments or strings, and so forth).

```

#include "eel.h"
#include "colcode.h"

int mymode_auto_show_delimiter()
{
    save_var point, case_fold = 1;
    save_var matchstart, matchend, abort_searching = 0;
    init_auto_show_delimiter(); // Must do this first.
    point -= parse_string(-1, "[a-z0-9_]+");
    *highlight_area_start[0] = point;
    if (parse_string(1, "</word>begin</word>")) {
        *highlight_area_end[0] = matchend;
        if (!re_search(1, "</word>end</word>"))
            return 2;
    } else if (parse_string(1, "</word>end</word>")) {
        *highlight_area_end[0] = matchend;
        if (!re_search(-1, "</word>begin</word>"))
            return 2;
    } else
        return 1;
    *highlight_area_start[1] = matchstart; // Mark the far end.
    *highlight_area_end[1] = matchend;
    modify_region(SHOW_MATCHING_REGION, MRTYPE, REGNORM);
    // Make the highlighting visible.
    return 3;
}

```

Figure 10.3: Highlighting keyword delimiters.

A language mode may also want to set things up so typing a closing delimiter momentarily moves the cursor back to show its matching pair. Binding keys like] and) to the command `show-matching-delimiter` will accomplish this.

DISPLAYING THE CURRENT FUNCTION'S NAME.

```
c_func_name_finder() // Sample C mode func name finder.
char display_func_name[];
char must_find_func_name;
int start_of_function;
set_display_func_name()
get_func_name(int idle)
```

A language mode may want to arrange for the name of the current function or similar to appear in the mode line, subject to the `display-definition` variable.

To do this, it must define a function named `modename_func_name_finder`, where *modename* is the mode's name as recorded in the `major_mode` variable. The function should write the current function's name to the `display_func_name` variable and return 1, also setting the `start_of_function` variable to a buffer position representing the start of the function, if it can. If point is not in a function, set `display_func_name` to an empty string and return 1.

Epsilon normally runs this function during idle time. If the user presses a key during this function, and the `must_find_func_name` variable is zero, it should stop any slow parsing if it can and return 0.

There are two subroutines that take advantage of such functions. The `set_display_func_name()` subroutine is what Epsilon calls when idle, to update the displayed function name.

The `get_func_name()` subroutine allows EEL code to take advantage of a mode's function name finder at other times. Pass a nonzero value for `idle` if you want it to give up and return should the user press a key. It returns 0 if Epsilon's idea of the function name in `display_func_name` was already up to date, 1 if it wasn't, but it now is, 2 if it couldn't be computed, or 3 if the function gave up to handle a waiting key. Both these functions set `start_of_function` to the start of the current function in this buffer if they can, or -1 otherwise.

HELPFUL SUBROUTINES.

Some subroutines help with mode-specific tasks.

```
int call_by_suffix(char *file, char *pattern)
int get_mode_variable(char *pat)
char *get_mode_string_variable(char *pat)
int get_mode_based_index(char *pat)
```

The `call_by_suffix()` subroutine constructs a function name based on the extension of a given file (typically the file associated with the current buffer). It takes the file name, and a function name with `%s` where the extension (without its leading ".") should be. For example, `call_by_suffix("file.cpp", "tag-suffix-%s")` looks for a subroutine named `tag-suffix-cpp`. (If the given file has no extension, the subroutine pretends the extension was "none".)

If there's no subroutine with the appropriate name, `call_by_suffix()` then replaces the `%s` with "default" and tries to call that function instead. The `call_by_suffix()` subroutine returns 1 if it found some function to call, or 0 if it couldn't locate any suitable function.

The `get_mode_variable()` subroutine searches for a function or variable with a name based on the current mode. Its parameter `pat` must be a printf-style format string, with a `%s` where the current mode's

name should appear. The subroutine will look for a function or variable with the resulting name. A variable by that name must be numeric; the subroutine will return its value. A function by that name must take no parameters and return a number; this subroutine will call it and return its value. In either case it will set the `got_bad_number` variable to zero. If `get_mode_variable()` can't locate a suitable function or variable, it sets `got_bad_number` nonzero.

The `get_mode_string_variable()` subroutine retrieves the value of a string variable whose name depends on the current mode. The name may also refer to a function; its value will be returned. It constructs the name by using `sprintf()`; `pat` should contain a `%s` and no other `%` characters; the current mode's name will replace the `%s`. If there's no such variable or function with that name, it returns `NULL`. The subroutine sets the `got_bad_number` variable nonzero to indicate that there was no such name, or zero otherwise.

The `get_mode_based_index()` subroutine looks for a name table entry of any sort (a function, variable, key table, etc.) with a name built by replacing the `%s` sequence in the specified pattern with the name of the current mode. If there is none, it substitutes "default" for the mode name and tries again. It returns the name table index of the entry it found, or zero if none.

```
int guess_mode_without_extension(char *res, char *pat)
```

The `guess_mode_without_extension()` subroutine tries to determine the correct mode for a file without an extension, or with an extension Epsilon doesn't recognize, by examining its text or complete file name. It can detect some Perl and C++ header files that lack any extension, as well as makefiles (based simply on the file's name) and various other sorts of files. If it can determine the mode, it uses `pat` as a pattern for `sprintf()` (so it should contain one `%s` and no other `%s`) and sets `res` to the `pat`, with its `%s` replaced by the mode name. Then it returns 1. If it can't guess the mode it returns 0.

```
mode_default_settings()
```

The `mode_default_settings()` subroutine resets a number of mode-specific variables to default settings. A command that establishes a mode can call this subroutine, if it doesn't want to provide explicit settings for all the usual mode-specific variables, such as comment pattern variables.

```
zeroed buffer (*buffer_maybe_break_line)();
int example_maybe_break_line(int type)
int generic_maybe_break_line(int type)
zeroed buffer int (*mode_restrict_break)();
int example_mode_restrict_break(int pos)
```

The auto-fill minor mode normally calls a function named `maybe_break_this_line()` to break lines. A major mode may set the buffer-specific function pointer `buffer_maybe_break_line` to point to a different function; then auto-fill mode will call that function instead, for possibly breaking lines as well as for turning auto-fill on or off, or testing its state.

A `buffer_maybe_break_line` function will be called with one numeric parameter. If 0 or 1, it's being told to turn auto-fill off or on. The function may interpret this request to apply only to the current buffer, or to all buffers in that mode. It should return 0.

If its parameter is 2, it's being asked whether auto-fill mode is on. It should return a nonzero value to indicate that auto-fill mode is on.

If its parameter is 3, it's being asked to perform an auto-fill, if appropriate, triggered by the key in the variable `key`, which has not yet been inserted in the buffer. It may simply return 1 if the line is not wide enough yet, or after it has broken the line. Epsilon will then insert the key that triggered the filling request. If it returns zero, Epsilon will skip inserting the key that triggered the filling.

Many language modes set `buffer_maybe_break_line` to point to the `generic_maybe_break_line()` function, which breaks within comments by using variables like `comment-start`, and doesn't break long lines outside comments. It works in languages that use simple one-line comments.

Even if a mode uses the standard `maybe_break_this_line()` subroutine to handle its line breaking, it can still limit where breaks may occur by setting the buffer-specific function pointer `mode_restrict_break` to point to a restriction function. A restriction function takes a parameter specifying the position of a space or tab character in the current buffer, and returns 1 if it's OK to break a line at that position, or 0 if it's not. The function must preserve `matchstart`, `matchend`, and `point`. In buffers where `mode_restrict_break` is zero, any space or tab character is a valid breaking position.

10.7.1 Language-specific Subroutines

```
int find_c_func_info(char *type, char *class,
                    char *func, int stop_on_key)
```

The `find_c_func_info()` subroutine gets info on the function or class defined at point in the current C-mode buffer, by parsing the buffer. It sets `class` to the class name of the current item, if any, and `func` to the function name if any. It sets `type` to "class", "struct", or "union" if it can determine which is appropriate. Outside a function or class definition, the above will be set to "". You may pass NULL for any of the above parameters if you don't need that information.

If `stop_on_key` is nonzero, and the user presses a key while the function is running, the function will immediately return -1 without setting the above variables. Otherwise the function returns a bit pattern: `CF_INFO_TYPE` if `type` was set non-empty; `CF_INFO_CLASS` if `class` was set non-empty; and `CF_INFO_FUNC` if `func` was set non-empty. In addition to zero, only these combination can occur:

CF_INFO_TYPE	CF_INFO_CLASS	CF_INFO_FUNC
•	•	•
	•	•
•	•	•

Chapter 11

Error Messages



This chapter lists some of the error messages Epsilon can produce, with explanations. In general, any error numbers produced with error messages are returned from the operating system.

Argument list mismatch in call. An EEL function was called with the wrong number of parameters. Perhaps you tried to call an EEL function by name, from the command line. Only functions that take no formal parameters can be called this way.

Can't find tutorial. Install first. Epsilon tried to load its tutorial file, since you started it with the `-teach` option, but can't find it. The tutorial is a file named `eteach`, located in Epsilon's main directory.

Can't interpret type of *variable-name*. You can only set or show variables that have numbers or characters in them.

COMSPEC missing from environment. Epsilon needs a valid COMSPEC environment variable in order to run another program. See page 11.

Couldn't exec: error *number*. You tried to run a program from within Epsilon, and Epsilon encountered an error trying to invoke that program. The *number* denotes the error code returned by the operating system. Also see the previous error.

Debug: can't read source file *filename*. Epsilon's EEL debugger tried to read an EEL source file, but couldn't find it. Epsilon gets a source file's pathname from the EEL compiler's command line. If you compiled an EEL file with the command `"eel dir/file.e"`, Epsilon will look for a file named `"dir/file.e"`. Check that your current directory is the same as when you ran the EEL compiler.

Don't know how to tag the file *filename*. Epsilon only knows how to tag files with certain extensions like `.c`, `.h`, `.e`, and `.asm`. Using EEL, you can tell Epsilon how to tag other types of files, though. See page 519.

Files not deleted. An error occurred when the `dired` command tried to delete the file or directory. You can only delete empty directories.

Invalid or outdated byte code file *filename*. The byte code file Epsilon tried to load was created with another version of Epsilon, was empty, or was illegal in some other way. Try compiling it again with the EEL compiler.

***filename* is not a directory.** You specified *filename* in an `-fs` flag, telling Epsilon to create its temporary files there, but it isn't a directory.

Macro definition buffer full: keyboard macro defined. You tried to define a macro of more than 500 keys from the keyboard. This might happen because you forgot to close a macro definition with the `Ctrl-x` command. If you really want to define such a big macro, use the command file mechanism (see page 164) or change the `MAX_MACRO` constant defined in `eel.h` and recompile `control.e` using EEL.


Macro nesting too deep. All macros canceled. An Epsilon keyboard macro can call another keyboard macro recursively (but only if the calling macro is defined by a command file—see page 156). To catch runaway recursive macros, Epsilon puts a limit on the depth of keyboard macro recursion. Epsilon allows unlimited *tail-recursion*: if a macro calls another macro with its last keystrokes, Epsilon finishes the original macro call before beginning the next one.

Only one window. The `diff` and `compare-windows` commands compare the current window with the next window on the screen, but there's only one window.

***function* undefined or of wrong type.** Epsilon initialized itself exclusively from a bytecode file (without reading a state file), since you gave the `-b` flag, but that file didn't define a function or variable that Epsilon needs to run. See page 552. To load a bytecode file, in addition to Epsilon's usual commands, use the `-l` flag, not the `-b` flag.

Appendix A

Index



- + command line option 10
- w EEL command line flag 398
- .bsc files for browsing 52
- .bsc files for tagging 52
- .directory.espell file 82
- .epsilon_vars file 129
- .espell files 82
- .mnu files 33
- .sbr files 53
- 8.3-format file names 514
- #messages# buffer 24
- #symbols# buffer 53

A

- abbreviate_file_name() subroutine 481, 511
- abort command 45, 106, **179**
- abort key 106
- ABORT_ASK textual macro 497
- ABORT_ERROR textual macro 446, 453, 497, 572
- abort_file_io primitive 265, 497
- abort_file_matching primitive 265, 509, 572
- ABORT_IGNORE textual macro 446, 497
- ABORT_JUMP textual macro 446, 453, 497, 572
- abort_key primitive 47, 265, 534
- abort_searching primitive 265, 446, 453
- about_box() primitive 575
- about-epsilon command 38, **179**
- absolute() primitive 510, 566
- add command line flag 14, 146
- add_buffer_when_displaying() subroutine 492
- add_buffer_when_idle() subroutine 556
- add_final_slash() primitive 511
- add_popup() primitive 463
- add_region() primitive 483
- add_tag() subroutine 519
- after-exiting color class 114
- after_loading() primitive 552
- align-by-tab command 81, **179**
- align-region command 81, **179**, 265
- align-region-extra-space variable 81, 265
- align-region-rules variable 81, 179, 265
- all_blanks() subroutine 452
- ALL_BORD() textual macro 463
- all-directory file variables 129
- all_must_build_mode primitive 266, 471
- all-read-only variable 266, 458
- alloc_spot() primitive 443
- allow_mouse_switching() subroutine 564
- already-made-backup buffer variable 266
- ALT() textual macro 557
- Alt-? key 37
- alt-invokes-menu variable 175, 266
- alt-numpad-keys variable 155, 266
- alt-prefix command 158, **179**
- alter_color() primitive 494
- anon-ftp-password variable 136, 266
- anonymous ftp 136
- another_process() primitive 531
- ansi-to-oem command 124, 179
- any_uppercase() subroutine 537
- API help 102
- append-next-kill command 61, **180**
- Application Data directory 13
- apply_defaults() primitive 552
- apropos command 37, 38, **180**
- argc primitive 266, 551
- argument command 155, **180**
- argument, numeric 27, 155
- argv primitive 551
- arrow keys 42
- ASCII characters 172
- aside() primitive 478
- ask_3way() subroutine 576
- ask_find_it() subroutine 496
- ask_key() subroutine 582
- ask_line_translate() subroutine 500
- ask_save_buffer() subroutine 498
- ask_yn() subroutine 575
- Asm mode 85
- asm-mode command 85, **180**
- aspell program 83
- assemble_mode_line() subroutine 470
- assigning to variables 159
- associations, file 146
- associativity 417
- ATTR_DIRECTORY textual macro 506
- ATTR_READONLY textual macro 506
- attr_to_rgb() primitive 494
- auto-fill-comment-rules buffer variable 105, 267
- auto-fill-indents buffer variable 77, 79, 128, 267
- auto-fill-mode buffer variable 128
- auto-fill-mode command 26, 77, 78, **180**
- auto-indent buffer variable 78, 267, 477
- auto-menu-bar variable 267
- auto-read-changed-file buffer variable 120, 267

auto-save-biggest-file variable 122, 267
 auto-save-count variable 122, 267
 auto-save-idle-seconds variable 122, 267
 auto-save-name variable 122, 268
 auto-save-tags variable 52, 268
 auto-show-adjacent-delimiter variable 44, 268
 auto-show-batch-delimiters variable 85, 268
 auto-show-c-delimiters variable 89, 268
 auto-show-conf-delimiters variable 90, 269
 auto-show-css-delimiters variable 92, 269
 auto-show-delimiter-delay variable 269
 auto-show-gams-delimiters variable 91, 269
 auto-show-html-delimiters variable 92, 269
 auto-show-matching-characters buffer variable 269, 585
 auto-show-nsis-delimiters variable 93, 269
 auto-show-perl-delimiters variable 94, 270
 auto-show-php-delimiters variable 95, 270
 auto-show-postscript-delimiters variable 95, 270
 auto-show-python-delimiters variable 96, 270
 auto-show-rmode-delimiters variable 96, 270
 auto-show-shell-delimiters variable 96, 270
 auto-show-tcl-delimiters variable 97, 270
 auto-show-tex-delimiters variable 97, 271
 auto-show-vbasic-delimiters variable 100, 271
 auto-show-vhdl-delimiters variable 100, 271
 autoload() primitive 550
 autoload_commands() primitive 549, 550
 autosaving files 121
 auxiliary files 12
 avoid-bottom-lines variable 115, 271, 462
 avoid-top-lines variable 115, 271, 462

B

-b command line flag 14, 397
 b_match() subroutine 568
 back-to-tab-stop command 80, **181**
 backup files 121
 backup-by-renaming variable 121, 271
 backup-name variable 121, 272
 backward-character command 42, **181**
 backward-delete-character command 58, **181**
 backward-delete-word command **181**
 backward-ifdef command 90, **181**
 backward-kill-level command 44, **181**
 backward-kill-word command 43, **181**

backward-level command 44, **181**
 backward-paragraph command 44, **181**
 backward-sentence command 43, **182**
 backward-word command 43, **182**
 Bash shell for Windows 151
 basic multilingual plane 137
 basic types 403
 Batch mode 85
 batch-auto-show-delim-chars variable 85, 272
 batch-mode command 85, **182**
 BBLANK textual macro 463
 BBOTTOM textual macro 463
 BC textual macro 474
 BDOUBLE textual macro 463
 BDOUBLEWIDTH textual macro 474
 beep-duration variable 116, 272, 522
 beep-frequency variable 116, 272, 522
 beginning-of-line command 42, **182**
 beginning-of-window command 107, **182**
 bell, setting 116
 bell-on-abort variable 117, 272
 bell-on-autosave-error variable 117, 272
 bell-on-bad-key variable 117, 272
 bell-on-completion variable 117, 273
 bell-on-date-warning variable 117, 273
 bell-on-read-error variable 117, 273
 bell-on-search variable 117, 273
 bell-on-write-error variable 117, 273
 BF_UNICODE textual macro 501
 BHEX textual macro 474
 binary constants 401
 binary files, editing 122
 Binary, in mode line 123
 bind-last-macro command 156, 157, **182**
 bind-to-key command 26, 157, 158, 175, **182**, 227
 in command file 165
 bind_universally() subroutine 583
 binding 26
 binding commands 157
 BLEFT textual macro 463
 block 416
 BM textual macro 474
 BMC textual macro 474
 BMP 137
 BNEWLINE textual macro 473, 474, 475
 BNONE textual macro 463
 BNORMAL textual macro 473, 474
 bold text 112

- bookmarks 50
- BORD() textual macro 463
- border-bottom variable 116, 273
- border-inside variable 116, 273
- border-left variable 116, 273
- border-right variable 116, 273
- border-top variable 116, 273
- BOTTOMRIGHT textual macro 461
- bprintf() primitive 440
- brace matching 44
- bracket matching 44
- break, eel keyword 414, 415
- break_into_numbers() subroutine 442
- Brief emulation 159
- brief-copy-region command **182**
- brief-cut-region command **182**
- brief-delete-region command **183**
- brief-delete-window command **183**
- brief-drop-bookmark command **183**
- brief-end-key command **183**
- brief-home-key command **183**
- brief-jump-to-bookmark command **183**
- brief-keyboard command 159, **183**
- brief-open-line command **183**
- brief-resize-window command **184**
- brief-split-window command **184**
- BRIGHT textual macro 463
- browse-current-symbol command 54, **184**
- browse-mode command **184**
- browse-set-filter command 54, **184**, 274
- browse-set-usedby-filter command 54, **185**, 274
- browse-symbol command 53, 54, **185**, 274
- browser files for tagging 52
- browser-file variable 53, 242, 274
- browser-filter variable 54, 274
- browser-filter-usedby variable 54, 274
- browser-options variable 54, 274
- browsing source code 52
- BSINGLE textual macro 463
- BTAB textual macro 473, 474
- BTOP textual macro 463
- buf-accessed buffer variable 274
- buf-accessed-clock variable 275
- buf_delete() primitive 456
- buf_exist() primitive 456
- buf_go_line() subroutine 451
- buf_grab_bytes() subroutine 441
- buf_in_window() primitive 503
- buf_list() primitive 459
- buf_match() primitive 571
- buf_pipe_text() primitive 533
- buf_position_to_line_number() subroutine 452
- buf_printf() primitive 440
- _buf_readonly buffer variable 457
- buf_set_character_color() subroutine 487
- buf_size() subroutine 457
- buf_sort_and_uniq() subroutine 455
- buf_stuff() primitive 440
- buf_xfer() subroutine 441
- buf_xfer_colors() subroutine 441, 487
- buf_zap() primitive 456
- bufed command 117, 143, 144, 154, **185**, 200
- bufed-column-width variable 144, 275
- bufed-grouping variable 144, 275, 369
- bufed-show-absolute-path variable 144, 275
- bufed-width variable 144, 275
- buffer 23
 - commands 117
 - keyword 403, 424, 430
 - startup 23
 - storage class 161
- buffer number 456
- buffer, eel keyword 161, 424
- buffer_display_characters buffer variable 474
- buffer_flags() primitive 501
- buffer_ftp_activity variable 517
- buffer-grep command 49, **185**
- buffer_list() primitive 459
- buffer_maybe_break_line buffer variable 588
- buffer-not-saveable buffer variable 275, 502
- buffer_on_modify buffer variable 457
- buffer_printf() primitive 440
- buffer_size() subroutine 457
- buffer_sort() primitive 453
- buffer-specific variables 161, 403, 548
- buffer-spell-mode command 82, 83, **186**
- buffer_to_clipboard() primitive 522
- buffer_unchanged() primitive 502
- buffer_url primitive 515
- buffers_identical() subroutine 454
- bufname primitive 276, 457
- bufnum primitive 276, 457
- bufnum_to_name() primitive 456
- build_filename() subroutine 513
- build_first primitive 276, 471, 474
- build_mode() subroutine 470

build_prompt() subroutine 570
 build_window() primitive 468
 BUNICODE textual macro 398
 button_dialog() primitive 575
 byte, eel keyword 403
 byte_extension primitive 276, 550
 bytecode files 14, 168
 bytes_to_chars() primitive 542

C

C++ mode 86
 c-access-spec-offset buffer variable 87, 276
 c-align-break-with-case buffer variable 87, 276
 c-align-contin-lines buffer variable 88, 276
 c-align-contin-max-offset buffer variable 277
 c-align-contin-max-width buffer variable 277
 c-align-extra-space buffer variable 88, 277
 c-align-inherit buffer variable 87, 277
 c-align-open-paren buffer variable 88, 277
 c-align-selectors buffer variable 88, 277
 C_ALPHA textual macro 538
 c-auto-fill-mode buffer variable 89, 104, 277
 c-auto-show-delim-chars variable 278
 c-biggest-declarator buffer variable 278
 c-block-macro-close-pat variable 278
 c-block-macro-inner-pat variable 278
 c-block-macro-open-pat variable 278
 c-brace-offset buffer variable 87, 278
 c-case-offset buffer variable 87, 278
 c-close command 89, **186**
 C_CNTRL textual macro 538
 c-colon command 89, **186**
 c-color-preproc-if-pattern variable 89, 278
 c-contin-offset buffer variable 87, 278
 c-delete-trailing-spaces buffer variable 121, 278, 297
 C_DIGIT textual macro 538
 c-extra-keywords buffer variable 279
 c-fill-column buffer variable 89, 105, 279
 C_GRAPH textual macro 538
 c-hash-mark command 89, **186**
 c-ident color class 115
 c-indent buffer variable 87, 128, 279
 c-indent-after-extern-c buffer variable 87, 279
 c-indent-after-namespace buffer variable 87, 279
 c-label-indent buffer variable 87, 279
 c-look-back variable 279
 C_LOWER textual macro 538
 c-mode command 86, 89, **186**
 c_move_level() subroutine 452
 c-open command 89, **186**
 c-param-decl buffer variable 87, 279
 c-prototype-macro-names variable 279
 C_PUNCT textual macro 538
 c-quick-scan buffer variable 90, 280
 c-reindent-previous-line buffer variable 280, 297
 c-spell-options variable 280
 c-tab-always-indents buffer variable 86, 280
 c-tab-override variable 87, 280
 c-tagging-class variable 280
 c-top-braces buffer variable 87, 280
 c-top-contin buffer variable 87, 280
 c-top-struct buffer variable 87, 281
 c-top-template buffer variable 281
 C_UPPER textual macro 538
 C_XDIGIT textual macro 538
 call_by_suffix() subroutine 587
 call_dll() primitive 528
 call_mode() subroutine 496
 call_on_modify primitive 281, 457
 call_with_arg_list() primitive 544
 can-get-process-directory variable 281
 canceling a command 106
 capitalize-word command 63, 64, **186**
 capture-output variable 150, 281
 caret 111
 carriage return translation 122, 498
 case replacement 65
 case, changing 63
 case, eel keyword 413, 415
 case, of file names 126
 case-fold buffer variable 45, 75, 257, 281, 446, 539
 cast, function pointer 543
 catch_mouse primitive 281, 559
 CAUTIOUS textual macro 568
 cd command 118, 119, **187**, 286
 center-line command 80, **187**
 center-window command 107, **187**
 CF_INFO_CLASS textual macro 589
 CF_INFO_FUNC textual macro 589
 CF_INFO_TYPE textual macro 589
 change_buffer_name() primitive 456
 change-code-coloring command 115, **187**
 change-file-read-only command 120, **187**
 change-font-size command 113, **187**

- change-line-wrapping command 108, **188**
- change-modified command 121, **188**
- change-name command 162, **188**
 - in command file 167
- change-read-only command 119, 120, **188**
- change-show-spaces command 111, 112, **188**
- changed files, detecting 120
- char, eel keyword 403
- char_avail() primitive 556
- character class 68
- character constant 401
- character sets, converting 124
- character() primitive 440
- charncmp() primitive 539
- chars_to_bytes() primitive 542
- chdir() primitive 508, 510
- check_abort() primitive 534, 578
- check_dates() subroutine 506
- CHECK_DEVICE textual macro 505
- CHECK_DIR textual macro 505
- CHECK_FILE textual macro 505
- check_file() primitive 505
- check_modify() primitive 458
- CHECK_OTHER textual macro 505
- CHECK_PATTERN textual macro 505
- CHECK_PIPE textual macro 505
- CHECK_URL textual macro 505
- checking spelling 81
- chm files 102
- clean-customizations command **188**
- clean_mode() subroutine 471
- clear-process-buffer variable 151, 281
- clear-tags command 52, **189**
- CLIP_ADD_FORMAT textual macro 522
- CLIP_CONVERT_NEWLINES textual macro 522
- clip_mouse() subroutine 561
- clipboard, accessing the 61
- clipboard-access variable 61, 262, 282
- clipboard_available() primitive 522
- clipboard-convert-mac-lines variable 61, 282
- clipboard-convert-unicode variable 61, 282
- clipboard-format variable 62, 282
- clipboard_to_buffer() primitive 522
- closeback buffer variable 86, 283
- CMD_INDEX_KEY textual macro 557
- cmd_len primitive 283, 579
- cmd-line-session-file variable 283
- CMDCONCURSHELLFLAGS configuration variable 151
- CMDSHELLFLAGS configuration variable 150
- code coloring 114
- coding file variable 128
- col_search() subroutine 450
- ColdFusion elements 283
- coldfusion-empty-elements variable 283
- color class 113, 492, 495
- color class assertions 74
- color scheme 113
- color_c_from_here() subroutine 490
- color_c_range() subroutine 489
- color_class, eel keyword 425
- COLOR_DO_COLORING textual macro 284
- color-html-look-back variable 283
- COLOR_IGNORE_INDENT textual macro 284, 490
- COLOR_IN_PROGRESS textual macro 284
- COLOR_INVALIDATE_BACKWARD textual macro 284, 490
- COLOR_INVALIDATE_FORWARD textual macro 284, 490
- COLOR_INVALIDATE_RESETS textual macro 284, 490
- color-look-back variable 115, 283, 490
- COLOR_MINIMAL textual macro 284
- color-names variable 284
- COLOR_RETAIN_NARROWING textual macro 284, 490
- color_scheme, eel keyword 426
- COLOR_STRIP_ATTR() textual macro 493
- color-whole-buffer variable 115, 284
- coloring-flags buffer variable 284, 490
- colors, changing 113
- column editing 63
- column number, always displaying 331
- column_color_searching() subroutine 450
- column_in_window primitive 284, 468
- column_to_pos() subroutine 475
- columnize_buffer_text() subroutine 454
- comm_dlg_color() primitive 575
- command
 - defined 404
 - eel keyword 404, 430
- command file
 - bind-to-key 165
 - change-name 167
 - create-prefix-command 166
 - define-macro 166
- command files 162, 164
- command history 30
- command line

- for EEL 397
- for Epsilon 9
- command, eel keyword 404
- comment-begin buffer variable 104, 284, 285, 585
- comment-column buffer variable 104, 284
- comment-end buffer variable 104, 284, 285, 585
- comment-pattern buffer variable 104, 285, 585
- comment-region command 105, **189**, 285, 350
- comment-region-options variable 104, 285
- comment-repeat-indentation-lines buffer variable 285
- comment-start buffer variable 104, 285, 585
- commenting commands 104
- comments in EEL 401
- common_file_dlg() primitive 574
- common-open-curdir variable 285, 574
- common-open-use-directory variable 125, 286
- COMP_FILE textual macro 568
- COMP_FOLD textual macro 568
- comp_read() subroutine 568
- comp_tab primitive 570
- compare_buffer_text() primitive 454
- compare_chars() primitive 539
- compare_dates() subroutine 506
- compare-files command 56, 57, **189**
- compare-files-ignore-spaces variable 56, 286
- compare-files-method variable 56, 286
- compare-files-mode command **189**
- compare-sorted-windows command 76, **189**
- compare-to-prior-version command 55, 56, **189**, 286
- compare-to-prior-version-style variable 56, 286
- compare-windows command 54, 56, **190**
- compare-windows-ignores-space variable 54, 286
- compile-asm-cmd variable 85, 287
- compile-buffer command 154, 155, **190**, 290
- compile-buffer-cmd buffer variable 287
- compile-c-cmd variable 287
- compile-c-cmd-unix variable 287
- compile-command file variable 154
- compile-cpp-cmd variable 154, 287
- compile-cpp-cmd-unix variable 287
- compile-csharp-cmd variable 287
- compile-eel-cmd variable 288
- compile-eel-dll-flags variable 154, 288
- compile-gams-cmd variable 288
- compile-html-cmd variable 288
- compile-idl-cmd variable 288
- compile-in-separate-buffer variable 154, 288
- compile-java-cmd variable 288
- compile-latex-cmd variable 98, 288
- compile-makefile-cmd variable 93, 289
- compile-makefile-cmd-unix variable 289
- compile-perl-cmd variable 94, 289
- compile-php-cmd variable 289
- compile-python-cmd variable 96, 289
- compile-rmode-cmd variable 289
- compile-tcl-cmd variable 289
- compile-tex-cmd variable 98, 289
- compile-typescript-cmd variable 290
- compile-vbasic-cmd variable 290
- compile-vhdl-cmd variable 290
- compile-xml-cmd variable 290
- compiler help 102
- complete() subroutine 569
- completion 28
 - adding your own 568
- completion, excluding files 30, 126
- completion_column_marker variable 569
- completion_lister variable 569
- completion-pops-up variable 28, 290
- complex scripts 137
- compressed files 119
- COMSPEC environment variable 11, 150
- concur_activity() subroutine 531
- concur_shell() primitive 530
- concurrent process 150
- concurrent-compile buffer variable 154, 290
- concurrent-make variable 154, 290
- COND_KEY textual macro 535
- COND_PROC textual macro 535
- COND_PROC_EXIT textual macro 535
- COND_RETURN_ABORT textual macro 535
- COND_TRUE_KEY textual macro 535
- Conf mode 90
- conf-auto-show-delim-chars variable 90, 290
- conf-mode command 90, **190**
- configuration variable 10
 - CMDCONCURSHELLFLAGS 151
 - CMDSHELLFLAGS 150
 - EEL 397
 - EPSCOMSPEC 11, 150, 151
 - EPSCONCURCOMSPEC 151
 - EPSCONCURSHELL 151
 - EPSCUSTDIR 11
 - EPSILON 14, 551
 - EPSMIXEDCASEDRIVES 126

- EPSPATH 12, 145, 397, 514
- EPSSHELL 12, 150, 151
- ESESSION 146
- INTERCONCURSHELLFLAGS 151
- INTERSHELLFLAGS 150
- configure-epsilon command 53, 146, **190**
- console-ansi-font variable 291, 538
- constants 401
- context help 102
- context-help command 102, 104, **190**
- context-help-default-rule variable 103, 291
- context_help_man() subroutine 104
- context_help_perldoc() subroutine 104
- context-help-rule-asm-unix variable 291
- context-help-rule-asm-windows variable 291
- context-help-rule-c-unix variable 291
- context-help-rule-c-windows variable 291
- context-help-rule-eel variable 291
- context-help-rule-gams variable 291
- context-help-rule-html variable 291
- context-help-rule-java variable 292
- context-help-rule-jscript variable 292
- context-help-rule-latex variable 292
- context-help-rule-nsis variable 292
- context-help-rule-objc variable 292
- context-help-rule-perl variable 292
- context-help-rule-php variable 292
- context-help-rule-postscript variable 292
- context-help-rule-python variable 292
- context-help-rule-shell variable 292
- context-help-rule-tcl variable 293
- context-help-rule-tex variable 293
- context-help-rule-vbasic-unix variable 293
- context-help-rule-vbasic-windows variable 293
- context-help-rule-vbscript variable 293
- context-help-rule-vhdl variable 293
- context-help-rule-xml variable 293
- context-help-site-search variable 103, 293
- context-help-site-search-lucky variable 103, 293
- context_help_windows_compilers() subroutine 104
- context-menu command 175, **190**
- continue, eel keyword 414, 415
- control characters 110
- control chars, in searches 45
- CONV_TO_16 textual macro 501
- conversion of variables 416
- convert_to_8_3_filename() primitive 514
- converting encodings 137, 501
- copy_buffer_variables() primitive 548
- copy_expanding() subroutine 541
- copy-file-name command 81, **191**
- copy-formatting-as-html command 62, **191**
- copy-include-file-name command 81, 119, **191**, 522
- copy-include-file-name-batch variable 294
- copy-include-file-name-latex variable 294
- copy-include-file-name-options variable 294
- copy-include-file-name-perl variable 294
- copy-include-file-name-shell variable 294
- copy-include-file-name-tex variable 294
- copy_key_table() subroutine 580
- copy_line_to_clipboard() subroutine 81, 522
- copy-rectangle command 63, **191**
- copy-region command 60, **191**
- copy-to-clipboard command 61, 62, **191**
- copy-to-file command 122, **191**
- copy-to-scratch command 61, **192**
- copyfile() primitive 504
- copying files 141
- copying text 59
- copyright, Epsilon iii
- count-lines command 108, **192**
- count_lines_in_buf() subroutine 452
- count_matches_in_buffer() subroutine 450
- count_windows_with_buf() primitive 503
- count-words command 107, 108, **192**
- CPROP_CTYPE textual macro 538
- CPROP_DEF_REPL textual macro 538
- CPROP_DISP_WIDTH textual macro 538
- CPROP_FOLD textual macro 538
- CPROP_TOLOWER textual macro 538
- CPROP_Toupper textual macro 538
- create() primitive 456
- create_dired_listing() subroutine 509
- create-file-associations command 192
- create_invisible_window() primitive 526
- create-prefix-command command 157, 158, **192**
- in command file 166
- create-variable command 162, **192**
- CSS mode 91
- css-auto-show-delim-chars variable 294
- css-indent buffer variable 92, 294
- css-mode command 93, **192**
- CTRL() textual macro 557
- Ctrl_ 37

ctrl-prefix command 158, **192**
 CTRLIFY_KEY textual macro 582
 cua-keyboard command 159, **192**
 curchar() primitive 440
 curl 131
 current buffer 25
 current window 25
 current_column() primitive 475
 current-output-color buffer variable 294
 current_time_and_date() subroutine 527
 curses program 17
 cursor-blink-period variable 295
 cursor_shape primitive 295, 482
 CURSOR_SHAPE() textual macro 482
 cursor_to_column primitive 295, 476
 customization directory 13
 cx_tab variable 580
 Cygwin environment 10, 102
 cygwin-filenames variable 10, 295
 cygwin-root variable 296

D

-d command line flag 14, 397
 date-format variable 81, 217, 296
 -dde command line flag 16
 DDE messages, sending 524
 dde_close() primitive 524
 dde_execute() primitive 524
 dde_open() primitive 524
 DEBUG_COLORING textual macro 491
 debug-text color class 114
 debugger 169
 decimal constant 401
 declaration 404
 declarator 405
 Def, in mode line 155
 default color class 114
 default value 161, 403
 default, eel keyword 414, 415
 default-add-final-newline variable 121, 296
 default-character-set variable 297
 default-color-spell-word-pattern variable 83, 297
 default-delete-trailing-spaces variable 121, 297
 default_fold() subroutine 449
 default_move_level() subroutine 452

default-read-encoding variable 137, 297
 default-reindent-previous-line variable 297
 default_replace_string() subroutine 449
 default_search_string() subroutine 449
 default-spell-options variable 82, 297
 default-spell-word-pattern variable 83, 298
 default-state-file-name variable 298
 default-translation-type variable 123, 298, 499
 default_when_displaying() subroutine 492
 default-word variable 43, 298
 default-write-encoding variable 137, 298
 #define preprocessor command 51, 397, 398
 define_color_class() primitive 494
 define-macro, in command file 166
 defined(), eel keyword 400
 delay() primitive 535
 delayed_say() primitive 479
 delete vs. kill 59
 delete() primitive 440
 delete-blank-lines command 58, **193**
 delete_buffer() primitive 456
 delete_buffer_when_displaying() subroutine 492
 delete_buffer_when_idle() subroutine 556
 delete-character command 58, **193**
 delete-current-line command **193**
 delete_file() primitive 503
 delete-hacking-tabs buffer variable 58, 181, 298, 299, 350
 delete-horizontal-space command 58, **193**
 delete_if_highlighted() subroutine 440
 delete-matching-lines command 65, **193**
 delete-name command 115, 162, **193**, 257
 delete-options variable 58, 299
 delete-rectangle command 63, **193**, 219, 325
 delete-region command 60, 61, **193**
 delete-to-end-of-line command **193**
 delete_user_buffer() subroutine 457
 deleting commands or variables 162
 deleting files 141
 describe-command command 37, 38, **194**
 describe-key command 37, 38, **194**
 describe-variable command 38, **194**
 desktop icon, running Epsilon from a 148
 detect_dired_format() subroutine 510
 detect-encodings variable 137, 297, 299
 detecting changed files 120
 Developer Studio, integrating with 147
 device files, ignoring 48

- DI_LINEINPUT textual macro 570
- DI_SEARCH textual macro 570
- DI_VIEW textual macro 570
- DI_VIEWLAST textual macro 570
- diacritical marks 136
- dialog_checkboxes primitive 577
- dialog-regex-replace command 65, **194**
- dialog-replace command 65, **194**
- dialog-reverse-search command 47, 48, **194**
- dialog-search command 47, 48, **194**
- dictionary lookup 81
- diff command 54, 56, **194**
- diff-match-characters variable 299
- diff-match-characters-limit variable 299
- diff-match-lines variable 54, 300
- diff-mismatch-lines variable 54, 300
- diff-precise-limit variable 300
- ding() primitive 522
- dir command line flag 14
- directory name, avoid typing 118
- directory, setting current 118
- directory_flags primitive 300
- directory-wide file variables 129
- dired command 10, 130, 139, 140, **194**, 202, 304, 509, 517
 - and find-file 118
- dired-24-hour-time variable 300
- dired-buffer-pattern buffer variable 300
- dired-confirmation variable 300
- dired-format buffer variable 300, 510
- dired-groups-dirs variable 300
- dired-layout variable 142, 195, 301
- dired-live-link-limit variable 301
- dired-mode command **195**
- dired_one() subroutine 509
- dired-show-dotfiles variable 301, 310
- dired-sort command **197**
- dired-sorts-files variable 301
- dired_standardize() primitive 509
- disable_dialog_controls primitive 578
- discardable-buffer buffer variable 301, 502
- disk management 139
- DISPLAY environment variable 11
- display-buffer-info command **197**
- display-c-options variable 301
- _display_characters primitive 474
- _display_class primitive 472, 473, 474
- display-column window variable 108, 302
- display-definition variable 280, 302, 334, 587
- display_dialog_box() primitive 576
- display-func-name variable 302
- display_more_msg() subroutine 471
- display_scroll_bar primitive 302, 564
- display_width() primitive 475
- displaying special characters 110
- displaying variables 159
- divide_url() subroutine 518
- DLLs, under Windows 528
- do, eel keyword 413
- do_again() primitive 581
- do_buffer_sort() subroutine 453
- do_buffer_to_hex() primitive 454
- do-c-indent command 89, **197**
- do_color_searching() subroutine 74, 447, 448, 450
- do_command() primitive 579, 581
- do_compare_sorted() subroutine 455
- do_compile() subroutine 190
- do_dired() primitive 509
- do_display_dialog_box() subroutine 577
- do_drop_matching_lines() subroutine 450
- do_execute_eel() subroutine 167
- do_file_match() subroutine 571
- do_file_read() subroutine 496
- do_find() subroutine 496
- do_ftp_op() subroutine 516
- do_insert_file() subroutine 503
- do_push() subroutine 530
- do_readonly_warning() subroutine 451, 496
- do_recursion() primitive 535
- do_remote_dired() subroutine 509
- do_resume_client() primitive 523
- do_save_file() subroutine 498
- do_save_state() subroutine 550
- do_searching() subroutine 448
- do_set_mark() subroutine 444
- do_shift_selects() subroutine 486
- do_sort_region() subroutine 453
- do_telnet() subroutine 515, 516
- do_topkey() primitive 578, 581
- do_uniq() subroutine 455
- do-vbasic-indent command **197**
- do_when_exiting_ subroutines 535
- do_when_idle_ subroutines 556
- do_when_make_mode_ subroutines 472
- do_when_repeating_ subroutines 556
- documentation, online 38

Documents and Settings directory 13
 _doing_input primitive 570
 DOS, in mode line 123
 DOS-format file names 514
 double_click_time primitive 302, 560
 down-line command 26, 42, **198**
 download_file_to_disk() subroutine 515
 drag_drop_handler() subroutine 523
 drag_drop_result() primitive 523
 dragging text 31
 draw-column-markers variable 111, 302
 draw-focus-rectangle variable 111, 303
 draw-line-numbers buffer variable 107, 111, 138, 303, 326, 331, 345, 365
 drive-match-types variable 139, 303
 drop_all_colored_regions() subroutine 492
 drop_buffer() subroutine 457
 drop_coloring() subroutine 492
 drop_dots() subroutine 509
 drop_final_slash() primitive 511
 drop_name() primitive 544, 545
 drop_pending_says() primitive 478
 DSABORT textual macro 448
 DSBAD textual macro 448
 DVI files, previewing 98, 218
 dvi viewers 98
 dynamic-link libraries, under Windows 528

E

-e command line flag 397
 early_init() subroutine 552
 EBADENCODE textual macro 501
 echo area 24
 _echo_display_class variable 474
 echo-line variable 115, 303
 ECOLOR_COPY textual macro 493
 ECOLOR_UNKNOWN textual macro 493
 edit-customizations command 13, 164, **198**
 edit-variables command 161, 162, **198**, 336
 edoc file 14, 38
 EEL 168
 EEL configuration variable 397
 EEL textual macro 398
 eel_compile() primitive 549
 EEL_PTR, type definition 546
 eel-tab-override variable 87, 303
 eel-version variable 303

EFONT_BOLD textual macro 427, 493
 EFONT_ITALIC textual macro 427, 493
 EFONT_UNDERLINED textual macro 427, 493
 eight bit characters 474
 einit-file-name variable 164, 304
 einit.ecm file 15, 162, 164
 #elif preprocessor command 400
 #else preprocessor command 400
 EMACS 27
 encoding_from_name() primitive 500
 encoding_to_name() primitive 500
 encodings
 converting 137, 501
 end-kbd-macro command 157, **198**, 227
 end-of-line command 42, **198**
 end-of-window command 107, **198**
 end_print_job() primitive 526
 #endif preprocessor command 400
 enlarge-window command 110, **198**
 enlarge-window-horizontally command 110, **198**
 enlarge-window-interactively command 110, **199**
 enter-key command 77, 78, **199**
 environment variable
 reading 520
 environment variable COMSPEC 11, 150
 environment variable DISPLAY 11
 environment variable EPSRUNS 12
 environment variable HOME 295
 environment variable MIXEDCASEDRIVES 126, 127
 environment variable PATH 12
 environment variable SHELL 12, 150
 environment variable TEMP 12, 14
 environment variable TMP 14
 environment variable USE_DEFAULT_COLORS 113
 EPSCOMSPEC configuration variable 11, 150, 151
 EPSCONCURCOMSPEC configuration variable 151
 EPSCONCURSHELL configuration variable 151
 EPSCUSTDIR configuration variable 11
 EPSILON configuration variable 14, 551
 Epsilon Extension Language 168
 Epsilon, command 9
 epsilon-help-format-unix-gui variable 38, 304
 epsilon-help-format-win-console variable 38, 304
 epsilon-help-format-win-gui variable 38, 304
 epsilon-html-look-up command **199**
 epsilon-info-look-up command 37, 39, **199**
 epsilon-keyboard command 159, **199**

epsilon-manual command 39, **199**
 epsilon-manual-html command **199**
 epsilon-manual-info command 37, 39, **199**
 epsilon-manual-port variable 304
 epsilon-viewer script 142
 epsilon-xfer-helper file 135
 EPSMIXEDCASEDRIVES configuration variable 126
 EPSPATH configuration variable 12, 145, 397, 514
 EPSRUNS environment variable 12
 EPSSHELL configuration variable 12, 150, 151
 EREADABORT textual macro 265, 497, 572
 err_file_read() subroutine 497
 errno primitive 304, 503, 508, 572
 error() primitive 534, 537
 error_if_input() subroutine 465
 ERROR_PATTERN textual macro 153
 ESESSION configuration variable 146
 eshell file 14
 espell.lst and espell.srt files 82
 eswap file 14
 ETOOBIG textual macro 501
 ETRANSPARENT textual macro 427, 493
 eval command 176, **199**, 200
 evaluate_numeric_expression() subroutine 573
 EWRITEABORT textual macro 265, 497
 EXACTONLY textual macro 568
 exchange-point-and-mark command 60, **200**
 executable files, editing 122
 execute-eel command 176, 199, **200**
 execution profiler 391
 exist() primitive 456
 exit command 144, **200**, 536
 exit-level command 144, **200**, 535
 exit-process command 152, **200**, 348
 expand_display() primitive 474
 expand_string_template() subroutine 513, 541
 expand-wildcards variable 10, 304
 expire_message variable 478
 explicit-session-file variable 305
 Explorer, opening folder in 140
 export-colors command 113, 114, **200**
 expressions in EEL 418
 exptoi() subroutine 573
 EXTEND_SEL_KEY textual macro 559
 extended file patterns 138
 extension language 391
 extensions vs. macros 391
 extensions, file 84

extract_rectangle() subroutine 485

F

-f command line flag 397
 F1 key 37
 fallback-remote-translation-type variable 123, 305
 fallback_translation_type buffer variable 500
 far-pause variable 44, 202, 305
 -fd command line flag 14
 field names 407
 file
 edoc 38
 eshell 14
 readme.txt 19
 startup 162
 file associations 146
 file associations, opening via 140
 file dates 120
 file name patterns 138
 file name prompts 125
 file name template 122, 513
 file names, capitalization of 126
 file types, customizing list of 125
 file variables 127
 FILE_CONVERT_ASK textual macro 501
 FILE_CONVERT_QUIET textual macro 501
 FILE_CONVERT_READ textual macro 500
 file_convert_read() subroutine 496
 FILE_CONVERT_WRITE textual macro 501
 file_convert_write() subroutine 500
 file-date-skip-drives variable 121, 305
 file-date-tolerance variable 120, 305
 file_error() primitive 503
 file_info structure 505
 file-interpret-env-vars variable 30, 119, 125, 306
 FILE_IO_ATSTART textual macro 497
 file_io_converter variable 500
 FILE_IO_NEWFILE textual macro 497, 501
 FILE_IO_TEMPFILE textual macro 497
 file_match() primitive 571
 file-pattern-ignore-directories variable 139, 306
 file-pattern-rules variable 139, 306, 571
 file-pattern-unc-domains variable 139, 306
 file-pattern-wildcards variable 139, 307, 505
 file-query-replace command 65, **200**

- `file_read()` primitive 495
- `file-read-kibitz` variable 124, 307
- `file-variables-search-limit` variable 128, 307
- `file_write()` primitive 497, 551
- `file_write_newfile` variable 497, 501
- `filename` primitive 307, 502
- `filename_rules()` primitive 513
- `FILETYPE_AUTO` textual macro 495, 499
- `FILETYPE_BINARY` textual macro 498
- `FILETYPE_MAC` textual macro 498
- `FILETYPE_MSDOS` textual macro 498, 499
- `FILETYPE_UNIX` textual macro 498, 499
- `fill` column 77
- `Fill`, in mode line 77
- `fill-c-comment-plain` buffer variable 90, 307
- `fill-comment` command 90, **201**
- `fill-indented-paragraph` command 78, **201**
- `fill-mode` buffer variable 77, 308
- `fill-paragraph` command 78, **201**
- `fill_rectangle()` subroutine 485
- `fill-region` command 77, 78, **202**
- `filter-region` command 150, **202**
- `filter.txt` file 125
- filters, customizing 125
- `final_index()` primitive 543
- `final-macro-pause` variable 308
- `find_buffer_prefix()` subroutine 570
- `find_c_func_info()` subroutine 589
- `find_data` variable 574
- `find-delimiter` command 44, **202**
- `find_dialog()` primitive 574
- `find_dialog_say()` primitive 574
- `find-file` command 51, 118, 119, 130, 141, 153, **202**, 496
 - and `dired` 118
- `find_group()` primitive 447
- `find_in_other_buf()` subroutine 496
- `find_index()` primitive 543
- `find_it()` subroutine 496, 585
- `find-lines-visible` variable 308
- `find-linked-file` command 119, **203**
- `find-linked-file-ignores-angles` variable 308
- `find_next_entry()` primitive 580
- `find-oem-file` command 124, **203**
- `find-read-only-file` command 120, **203**
- `find_remote_file()` subroutine 496
- `find-unconverted-file` command **203**
- `finger_user()` primitive 515
- `finish_up()` subroutine 552
- `first_window_refresh` primitive 308, 492
- `fix_cursor()` subroutine 482
- `fix_key_table()` subroutine 580
- `fix_region()` subroutine 484
- `fix_window_start()` subroutine 467
- `FKEY()` textual macro 557
- flags
 - for `EEL` 397
 - for `Epsilon` 13
- `FM_FOLD` textual macro 571
- `FM_NO_DIRS` textual macro 568, 571
- `FM_ONLY_DIRS` textual macro 568, 571
- `fnamecmp()` subroutine 513
- `FNAMLEN` textual macro 392
- `FNT_DIALOG` textual macro 482
- `FNT_PRINTER` textual macro 482
- `FNT_SCREEN` textual macro 482
- `FOLD` textual macro 448
- `follow-mode` command 109, **203**
- `follow-mode-on` buffer variable 308
- `follow-mode-overlap` variable 109, 203, 308
- font styles 112
- `font-dialog` variable 113, 308
- `font-fixed` variable 112, 309
- `font-printer` variable 113, 309
- `font-styles` variable 309
- `font-styles-tolerance` variable 113, 309
- fonts, setting 112
- `for`, `eel` keyword 413
- `force-common-file-dialog` command 125, **203**
- `FORCE_MODE_LINE` textual macro 470
- `force-remote-translation-type` variable 123, 298, 309
- `force-save-as` buffer variable 310
- `force_to_column()` subroutine 475
- foreign characters 136, 473
- format string 480
- `format_date()` subroutine 506
- `format_file_date()` subroutine 506
- `forward-character` command 42, **204**
- `forward-ifdef` command 90, **204**
- `forward-level` command 44, **204**
- `forward-paragraph` command 44, 181, **204**, 225
- `forward-search-again` command 47, 48, **204**
- `forward-sentence` command 43, **204**, 220
- `forward-word` command 43, **204**
- `forward-word-to-start` variable 310
- `FPAT_COMMA` textual macro 505

FPAT_CURLY_BRACE textual macro 505
 FPAT_FOLD textual macro 540
 FPAT_IGNORE_SQUARE_BRACKETS textual macro 540
 FPAT_PERMIT_NO_URLS textual macro 306
 FPAT_SEMICOLON textual macro 505
 FPAT_SKIP_DIR_SYMLINKS textual macro 306
 FPAT_SKIP_FILE_SYMLINKS textual macro 306
 FPAT_SKIP_RECUR_SYMLINKS textual macro 306
 FPAT_SQUARE_BRACKET textual macro 505
 fpatmatch() primitive 540
 free() primitive 542
 free_spot() primitive 443
 -fs command line flag 14, 543
 FSA_NEWFILE textual macro 310
 FSA_READONLY textual macro 310
 FSYS_CASE_IGNORED textual macro 513
 FSYS_CASE_MASK textual macro 513
 FSYS_CASE_PRESERVED textual macro 513
 FSYS_CASE_SENSITIVE textual macro 513
 FSYS_CASE_UNKNOWN textual macro 513
 FSYS_CDROM textual macro 303, 363, 513
 FSYS_LOCAL textual macro 303, 513
 FSYS_NETWORK textual macro 303, 363, 513
 FSYS_REMOVABLE textual macro 303, 363, 513
 FSYS_SHORT_NAMES textual macro 513
 FTP URL 130
 ftp_activity() subroutine 517
 FTP_ASCII textual macro 516
 ftp-ascii-transfers variable 123, 131, 310, 516
 ftp-compatible-dirs variable 131, 310, 516
 FTP_LIST textual macro 516
 FTP_MISC textual macro 516
 ftp_misc_operation() subroutine 516
 ftp_op() primitive 516, 518
 FTP_OP_MASK textual macro 516
 ftp-passive-transfers variable 131, 311
 FTP_PLAIN_LIST textual macro 516
 FTP_RECV textual macro 516
 FTP_SEND textual macro 516
 FTP_USE_CWD textual macro 516
 FTP_WAIT textual macro 516
 full_getkey() subroutine 582
 full_key primitive 311, 555
 full-path-on-mode-line variable 311
 full_redraw primitive 311, 471, 472
 function 422
 function keys 173
 function name, displaying 302

function, pointer to 543
 fundamental-auto-show-delim-chars variable 85, 311
 fundamental-mode command 85, **205**
 fundamental-spell-options variable 311
 fwd-search-key variable 47, 312

G

GAMS files 288
 GAMS mode 90
 gams-auto-show-delim-chars variable 312
 gams-files variable 90, 312
 gams-mode command 91, **205**
 general_matcher() primitive 570
 generic_key() primitive 555
 generic_maybe_break_line() subroutine 589
 -geometry command line flag 14
 get_any() subroutine 567
 get_background_color() primitive 493
 GET_BORD() textual macro 463
 get_buf() subroutine 566
 get_buf_modified() subroutine 502
 get_buf_point() subroutine 457
 get_buffer_directory() subroutine 508
 get_buffer_filename() subroutine 502
 get_character_color() primitive 488
 get_choice() subroutine 573
 get_cmd() subroutine 567
 get_color_scheme_variable() subroutine 493
 get_column() subroutine 475
 get_command_index() subroutine 567
 get_customization_directory() primitive 508
 get_default_translation_type() subroutine 496
 get_direction() subroutine 538
 get_dired_item() subroutine 510
 get_doc() subroutine 554
 GET_ENCODING() textual macro 499
 get_executable_directory() primitive 512
 get_executable_file() primitive 512
 get_extension() primitive 511
 get_fallback_translation_type() primitive 500
 get_file() subroutine 566
 get_file_dir() subroutine 566
 get_file_read_kibitz() primitive 502
 get_file_read_only() primitive 504
 get_foreground_color() primitive 493
 get_func() subroutine 567

- get_func_name() subroutine 587
- get_indentation() subroutine 475
- get_key_choice() subroutine 573
- get_key_response() subroutine 576
- get_keycode() primitive 558, 582
- get_line_from_buffer() subroutine 442
- GET_LINE_TRANSLATE() textual macro 499
- get_macname() subroutine 567
- get_macro() primitive 582
- get_mode_based_index() subroutine 588
- get_mode_string_variable() subroutine 588
- get_mode_variable() subroutine 587
- get_movement_or_release() subroutine 562
- get_num_var() primitive 545
- get_number() subroutine 572
- get_password() subroutine 518
- get_profile() primitive 554
- get_range() primitive 579
- get_search_string() subroutine 449
- get_spot() primitive 444
- get_str_auto_def() subroutine 572
- get_str_var() primitive 545
- get_strerror() subroutine 572
- get_string() subroutine 572
- get_strnone() subroutine 572
- get_strpopup() subroutine 572
- get_substitution_variables() primitive 520
- get_tagged_region() primitive 488
- get_tail() primitive 511
- get_tempfile_name() primitive 497
- get_url_file_part() subroutine 518
- get_var() subroutine 567
- get_wattrib() primitive 466
- get_window_info() subroutine 462
- get_window_pos() primitive 467
- GETBLUE() textual macro 427
- getc() primitive 508
- getenv() primitive 520
- GETFOCUS textual macro 565
- GETGREEN() textual macro 427
- gethostname() primitive 517
- getkey() subroutine 556, 578, 581
- GETRED() textual macro 427
- give_begin_line() subroutine 451
- give_end_line() subroutine 451
- give_line_translate() subroutine 499
- give_position() subroutine 453
- give_position_at_column() subroutine 475
- give_prev_buf() subroutine 467
- give_window_space() primitive 459
- glibc 5
- global variable 402
- global-spell-options variable 82, 312
- go_line() subroutine 451
- goal-column buffer variable 312
- got-bad-number variable 312, 573
- goto, eel keyword 415, 416
- goto-beginning command 42, **205**
- goto-end command 42, **205**
- goto-line command 107, 108, **205**
- goto-next-declaration command 101, **205**
- goto-next-definition command 101, **205**
- goto-previous-declaration command 101, **205**
- goto-previous-definition command 101, **206**
- goto-tag command 50, 52, **206**
- goto_url file 41
- grab() primitive 441
- grab_buffer() subroutine 441
- grab_expanding() subroutine 441
- grab_full_line() subroutine 441
- grab_line() subroutine 441
- grab_line_offset() subroutine 442
- grab_numbers() subroutine 442
- grab_string() subroutine 442
- grab_string_expanding() subroutine 442
- graphics characters 110, 136
- grep command 49, **206**, 314
- grep-default-directory variable 48, 312, 314
- grep-empties-buffer variable 49, 313
- grep-ignore-file-basename variable 48, 242, 313
- grep-ignore-file-extensions variable 48, 242, 313
- grep-ignore-file-max-size variable 48, 242, 313
- grep-ignore-file-pattern variable 48, 242, 313
- grep-ignore-file-types variable 49, 242, 313
- grep-include-timestamp variable 313
- grep-keeps-files variable 49, 313
- grep-mode command **207**
- grep-on-changed-file variable 314
- grep-prompt-with-buffer-directory variable 48, 314
- grep-show-absolute-path variable 314
- GREYBACK textual macro 557
- GREYENTER textual macro 557
- GREYEQUAL textual macro 557
- GREYESC textual macro 557

GREYHELP textual macro 557
 GREYMINUS textual macro 557
 GREYPLUS textual macro 557
 GREYSLASH textual macro 557
 GREYSTAR textual macro 557
 GREYTAB textual macro 557
 grouping of EEL operators 417
 guess_mode_without_extension() subroutine 588
 gui_cursor_shape primitive 314, 482
 GUI_CURSOR_SHAPE() textual macro 482
 gui-menu-file variable 33, 315

H

hack_tabs() subroutine 476
 halt_process() primitive 532
 has_arg primitive 315, 535, 578, 579
 has_feature primitive 521
 help command 37, 38, **207**
 file 14
 help, getting 37
 help_on_command() subroutine 554
 help_on_current() subroutine 554
 HELPKEY textual macro 576
 hex constants 401
 entering interactively 159
 hex display 110
 hex-mode command 84, **208**
 hex-overtyping-mode variable 315
 hex-tab-key command **208**
 _highlight_control primitive 484
 highlight_off() subroutine 484
 highlight_on() subroutine 484
 highlight-region command 60, **209**
 history of commands 30
 hlp files 102
 HOME environment variable 295
 hook
 when loading bytecode files 549
 when reading in a file 84
 when starting Epsilon 552
 horiz-border color class 114, 495
 horizontal scrolling 108
 HORIZONTAL textual macro 460
 horizontal() primitive 475
 host name, displaying 378
 host name, retrieving 517
 HTML mode 91

html-asp-coloring variable 92, 315
 html-auto-fill-combine variable 315
 html-auto-fill-mode variable 315
 html-auto-indent variable 91, 316
 html-auto-show-delim-chars variable 316
 html-backward-tag command 93, **209**
 html-close-last-tag command 93, **209**
 html-delete-tag command 93, **209**
 html-display-definition variable 316
 html-display-nesting-width variable 92, 316
 html-empty-elements variable 316
 html-fill-paragraph command **209**
 html-find-matching-tag command 93, **209**
 html-forward-tag command 93, **209**
 html-indent buffer variable 91, 316
 html-indent-cmd command **209**
 html-indenting-rules variable 91, 317
 html-list-element-nesting command 93, **209**
 html-list-mismatched-tags command 93, **210**
 html-mode command 92, **210**
 html_move_level() subroutine 452
 html-no-indent-elements variable 91, 317
 html-other-coloring variable 92, 318
 html-paragraph-elements variable 318
 html-paragraph-is-container buffer variable 91, 318
 html-php-coloring variable 92, 318
 html-prevent-coloring variable 91, 318
 html-recognize-coldfusion-comments buffer variable 319
 html-redirect-active-key command **210**
 html-reindent-previous-line variable 319
 html-spell-options variable 319
 html-style-rules buffer variable 319
 html-tag-match-look-back variable 319
 HtmlHelp files 102
 Http URL 131
 http_force_headers primitive 515
 http-force-headers variable 131
 http-log-request variable 131, 319
 http-proxy-exceptions variable 319
 http-proxy-port variable 319, 320
 http-proxy-server variable 319, 320
 http_retrieve() primitive 515
 HTTP_RETRIEVE_ONLY_HEADER textual macro 515
 HTTP_RETRIEVE_WAIT textual macro 515
 http-user-agent variable 320
 Https URL 131

I

- i command line flag 397
- IACT_AUTO_FILL textual macro 477
- IACT_AUTO_FILL_COMMENT textual macro 477
- IACT_AUTO_INDENT textual macro 477
- IACT_COMMENT textual macro 477
- IACT_FILL_COMMENT textual macro 477
- IACT_REINDENT_PREV textual macro 477
- identifiers 401
- IDL files 288
- idle-coloring-delay buffer variable 115, 320
- idle-coloring-size buffer variable 320
- #if preprocessor command 400
- if, eel keyword 277, 413
- ifdef lines, moving by 89
- #ifdef preprocessor command 400
- #ifndef preprocessor command 400
- ig_nore_file_extensions variable 513
- ignore-error variable 153, 320
- ignore-file-basename variable 320
- ignore-file-extensions variable 30, 126, 320
- ignore-file-pattern variable 320
- ignore-kbd-macro variable 321, 582
- ignore.lst file 82
- ignoring-file-change buffer variable 321
- import-colors command **210**
- import-customizations command 170, **210**
- in_bufed() subroutine 467
- in_echo_area primitive 321, 481, 578
- in_macro() primitive 556
- in-perl-buffer buffer variable 321
- in-shell-buffer buffer variable 321
- include preprocessor command 399
 - executed only once 430
- include-directories variable 119, 203, 321
- INCR textual macro 448
- incremental-search command 45, 48, **210**
- indent-comment-as-code variable 104, 211, 321
- indent-for-comment command 105, **211**
- indent_like_tab() subroutine 475
- indent-preprocessor-contin buffer variable 321
- indent-previous command 78, 80, **211**
- indent-region command 79, 80, **211**
- indent-rigidly command 79, 80, **211**
- indent_to_column() subroutine 475
- indent-to-tab-stop command 80, **212**
- indent-under command 79, 80, **212**
- indent-with-tabs buffer variable 63, 79, 321, 475
- indenter variable 477
- indenter_action primitive 477
- indenting 78
- indents-separate-paragraphs buffer variable 43, 204, 322
- index() primitive 539
- index_table() primitive 580
- info command 41, **212**
- info-backward-node command 40, **212**
- info-directory-node command 41, **212**
- info-follow-nearest-reference command 41, **212**
- info-follow-reference command 41, **213**
- info-forward-node command 40, **213**
- info-goto command 41, **213**
- info-goto-epsilon-command command 38, **213**
- info-goto-epsilon-key command 38, **213**
- info-goto-epsilon-variable command 38, **213**
- info-index command 41, **213**
- info-index-next command 41, **213**
- info-last command 41, **214**
- info-last-node command 41, **214**
- info-menu command 40, **214**
- info-mode command 41, **214**
- info-mouse-double command **215**
- info-next command 40, **215**
- info-next-page command 40, **215**
- info-next-reference command 41, **215**
- info-nth-menu-item command 41, **215**
- info-path-non-unix variable 40, 322
- info-path-unix variable 40, 322
- info-previous command 40, **215**
- info-previous-page command 40, **216**
- info-previous-reference command 41, **216**
- info-quit command 41, **216**
- info-recovering variable 322
- info-search command 41, **216**
- info-tagify command 40, 41, **216**
- info-top command 41, **216**
- info-up command 40, **216**
- info-validate command 40, 41, **216**
- Ini mode 93
- ini-mode command 93, **216**
- initial-tag-file variable 52, 322
- initialization
 - of Epsilon 14
 - of variables 410, 411
- inline_eel() subroutine 167
- insert() primitive 439

- insert-ascii command 57, 58, **217**
- insert-binding command 168, **217**
- insert-clipboard command 61, 62, **217**
- insert-date command 81, **217**
- insert-default-response variable 28, 60, 322
- insert-file command 119, **217**, 503
- insert-file-remembers-file variable 322
- insert-macro command 157, 168, **217**
- insert-scratch command 61, **217**
- insert_to_column() subroutine 475
- inserting characters 57
- installation 5
 - for DOS 9
 - for macOS 6
 - for OS/2 9
 - for Unix 5
- Installing Epsilon for DOS 9
- Installing Epsilon for macOS 6
- Installing Epsilon for OS/2 9
- Installing Epsilon for Unix 5
- int, eel keyword 403, 558
- int32, eel keyword 403
- integrate with Visual Studio 147
- integrating with Developer Studio 147
- IntelliMouse support 31, 380, 565
- INTERCONCURSHELLFLAGS configuration variable 151
- international characters 136, 473
- internationalization 136
- Internet 130
- INTERSHELLFLAGS configuration variable 150
- invisible_cmd() primitive 579
- invisible_window primitive 322, 466
- invoke_menu() primitive 525
- invoke-windows-menu command 175, 217
- invoking Epsilon 9
- IS_ALT_KEY() textual macro 557
- is_buf_in_window() subroutine 503
- is_buffer_in_window() subroutine 503
- IS_CTRL_KEY() textual macro 557
- is-current-window variable 323
- is_directory() primitive 504
- is_dired_buf() subroutine 509
- IS_EXT_ASCII_KEY() textual macro 561
- is_gui primitive 323, 520
- is_highlight_on() subroutine 484
- is_in_tree() subroutine 512
- is_key_repeating() primitive 556
- IS_LEADING_SURROGATE() textual macro 454
- IS_MOUSE_...() textual macros 561
- IS_MOUSE_CENTER() textual macro 561
- IS_MOUSE_DOUBLE() textual macro 561
- IS_MOUSE_DOWN() textual macro 561
- IS_MOUSE_KEY() textual macro 560
- IS_MOUSE_LEFT() textual macro 561
- IS_MOUSE_RIGHT() textual macro 561
- IS_MOUSE_SINGLE() textual macro 561
- IS_MOUSE_UP() textual macro 561
- is_mswindows primitive 521
- is-mswindows variable 323
- IS_MSWINDOWS_32BIT textual macro 323, 521
- IS_MSWINDOWS_64BIT textual macro 323, 521
- IS_NT textual macro 323, 520
- is_path_separator() primitive 511
- is_pattern() primitive 504
- is_process_buffer() primitive 531
- is_relative() primitive 510
- is_remote_buffer() subroutine 515
- is_remote_file() primitive 512
- IS_TRAILING_SURROGATE() textual macro 454
- IS_TRUE_KEY() textual macro 560
- is_unix primitive 520
- is-unix variable 323
- IS_UNIX_BSD textual macro 323, 520
- is_unix_flavor primitive 520
- is-unix-flavor variable 323
- IS_UNIX_LINUX textual macro 323, 520
- IS_UNIX_MACOS textual macro 323, 520
- IS_UNIX_TERM textual macro 323, 520
- IS_UNIX_XWIN textual macro 323, 520
- is_unsaved_buffer() subroutine 502
- IS_WIN_KEY() textual macro 561
- IS_WIN_PASSIVE_KEY() textual macro 561
- is_window() primitive 461
- isalnum() subroutine 537
- isalpha() primitive 537
- iscntrl() primitive 537
- isdigit() primitive 537
- isgraph() primitive 537
- isident() subroutine 537
- islower() primitive 537
- ISO 8859 character sets 136
- ispell program 83
- ISPOPOP textual macro 461
- isprint() primitive 537
- ISPROC_CONCUR textual macro 531

ISPROC_PIPE textual macro 531
 ispunct() primitive 537
 isspace() primitive 537
 ISTILED textual macro 461
 isupper() primitive 537
 isxdigit() primitive 537
 italic text 112
 iter primitive 323, 535, 572, 578, 579, 581

J

Java mode 86
 java-indent buffer variable 87, 323
 jump-to-column command 108, **218**
 jump-to-dvi command 98, 99, 100, **218**
 jump-to-dvi-command variable 99, 323
 jump-to-dvi-extension variable 99, 324
 jump-to-dvi-main-file variable 324
 jump-to-last-bookmark command 50, **218**
 jump-to-named-bookmark command 50, **218**

K

-ka command line flag 15
 keep-duplicate-lines command 76, **218**
 keep-matching-lines command 65, **218**
 keep-unique-lines command 76, **218**
 key primitive 324, 555, 581
 key table 424, 579
 key table, values in 543
 KEY_ALT textual macro 557
 key_binding variable 582
 key_code primitive 324, 559
 KEY_CTRL textual macro 557
 key-from-macro variable 324, 556
 key_is_button primitive 324, 566, 577
 KEY_PLAIN() textual macro 557
 key-repeat-rate variable 176, 324
 KEY_SHIFT textual macro 557
 key_t textual macro 558, 579
 key_type primitive 325, 559
 key_value() primitive 481, 558
 keyboard macro 155
 KEYDELETE textual macro 557
 KEYDOWN textual macro 557
 KEYEND textual macro 557
 KEYHOME textual macro 557
 KEYINSERT textual macro 557
 KEYLEFT textual macro 557

KEYPGDN textual macro 557
 KEYPGUP textual macro 557
 KEYRIGHT textual macro 557
 keys and commands 157
 Keystrokes and Commands: Bindings 26
 keystrokes, recording 155
 keytable 424, 579
 keytable, eel keyword 424, 430, 580, 584
 keytable, values in 543
 KEYUP textual macro 557
 keyword help 102
 kill buffers 59
 kill vs. delete 59
 kill-all-buffers command 117, 118, **219**
 kill-buffer command 118, **219**
 kill-buffers variable 60, 325
 kill-comment command 105, **219**
 kill-current-buffer command 118, **219**
 kill-current-line command 59, 61, **219**
 kill-level command 44, **219**
 kill-line command 60, **219**
 kill-process command 152, **219**
 kill-rectangle command 63, **219**, 325
 kill-rectangle-removes variable 325
 kill-region command 60, **220**
 kill-sentence command 43, **220**
 kill-to-end-of-line command 59, 61, **220**
 kill-window command 109, **220**
 kill-word command 43, **220**
 killing commands 59
 -ks command line flag 15
 KT_ACCENT textual macro 559
 KT_ACCENT_SEQ textual macro 559
 KT_EXTEND_SEL textual macro 559
 KT_MACRO textual macro 559
 KT_NONASCII textual macro 559
 KT_NONASCII_EXT textual macro 559
 KT_NORMAL textual macro 559

L

-l command line flag 15, 552
 language mode
 defining a new 584
 language modes 84
 last_index primitive 325, 554, 581
 last-kbd-macro command 156, 157, **220**
 last-show-spaces buffer variable 325

- last-window-color-scheme variable 325
 - LaTeX mode 97
 - latex-2e-or-3 variable 97, 325
 - latex-display-math-env-pat variable 325
 - latex-math-env-pat variable 325
 - latex-mode command 100, **220**
 - latex-non-text-argument variable 98, 326
 - latex-spell-ignore-pattern-prefix variable 326
 - latex-spell-options variable 326
 - latex-tag-keywords variable 98, 326
 - Latin 1 character set 136
 - lcs() primitive 455
 - lcs_char() primitive 455
 - leave() primitive 535, 536
 - leave_blank primitive 326, 552
 - leave_recursion() primitive 535, 578
 - _len_def_mac variable 556
 - level 44
 - libnss shared files 5
 - licensing, Epsilon iii
 - lifetime of variables 402
 - line number, displaying 107
 - line number, positioning by 107
 - line numbers, always displaying 331
 - #line preprocessor command 399
 - line scrolling 108
 - line translation 122, 498
 - line wrapping 108
 - line_in_window primitive 326, 468
 - line-number-width buffer variable 111, 303, 326
 - line_search() subroutine 450
 - line-to-bottom command 106, 107, **220**
 - line-to-top command 106, 107, **221**
 - lines_between() primitive 451
 - lisp commands 44
 - list-all command 168, 170, 171, 172, **221**
 - list_bindings() primitive 583
 - list-bookmarks command 50, **221**
 - list-changes command 171, 172, **221**
 - list-colors command **221**
 - list-customizations command 168, **221**, 387
 - list-debug command **221**
 - list-definitions command 90, 101, **222**
 - list-definitions-live-update variable 101, 326
 - list-files command 140, **222**
 - list_finder variable 569
 - list-make-preprocessor-conditionals command 93, **222**
 - list_matches() subroutine 569
 - list-preprocessor-conditionals command 90, **222**
 - list-undefined command 169, **222**
 - list-which-definitions variable 327
 - LISTMATCH textual macro 568
 - load-buffer command 157, 164, **222**
 - load-bytes command 169, **222**, 388
 - load-changes command 171, 172, **223**
 - load_commands() primitive 549
 - load-customizations variable 327
 - load_eel_from_path() subroutine 166, 549
 - load-fail-ok variable 327
 - load-file command 164, **223**
 - load_from_path() subroutine 549, 550
 - load_from_state primitive 327, 552
 - local variable 402
 - locate-file command 140, **223**
 - locate-path-unix variable 140, 223, 327
 - locate_window() subroutine 503
 - long lines 108
 - longjmp() primitive 536
 - look_file() subroutine 496
 - look_on_path() primitive 513
 - look_up_tree() subroutine 512
 - lookpath() primitive 513
 - LOSEFOCUS textual macro 565
 - low_file_match() subroutine 571
 - low-level operations 528
 - low_window_create() primitive 462
 - low_window_info() primitive 462
 - lowaccess() primitive 507
 - lowclose() primitive 507
 - lowercase-word command 63, 64, **223**
 - lowopen() primitive 507
 - lowread() primitive 507
 - lowseek() primitive 507
 - lowwrite() primitive 507
 - LR_BORD() textual macro 463
 - lugeps.ini file 148
 - lvalue expressions 418
- ## M
- m command line flag 15
 - Mac, in mode line 123
 - mac-framework-dirs variable 327
 - Macintosh files 122
 - Macintosh, running Epsilon on 6
 - macro-runs-immediately variable 327, 581

- macros vs. extensions 391
- macros, keyboard 155
- macros, types of 398
- mail-fill-paragraph command 78, **223**
- mail-quote-pattern variable 78, 223, 328
- mail-quote-region command 78, **223**
- mail-quote-skip variable 78, 223, 328
- mail-quote-text variable 78, 223, 328
- mail-unquote command 78, **223**
- main loop 578
- major modes 25
- major-mode buffer variable 328, 470
- make command 153, 154, 155, **224**, 290
- make utility program 388
- make_alt() subroutine 557
- make_anon_keytable() subroutine 580
- make_backup() primitive 504
- make_ctrl() subroutine 557
- make_dired() subroutine 509
- make_line_highlight() subroutine 486
- make_mode() subroutine 471
- make_pointer() primitive 529
- MAKE_RGB() textual macro 426
- make_title() primitive 470
- MAKE_TRANSLATE() textual macro 499
- makefile file 388
- Makefile mode 93
- makefile-mode command 93, **224**
- malloc() primitive 542
- man command 102, 104, **224**
- margin-right buffer variable 77, 328
- margins, setting printer 137
- mark 59
- mark primitive 328, 444
- mark-c-paragraph command **224**
- mark-inclusive-region command 63, **224**
- mark-line-region command 63, **224**
- mark-normal-region command 62, **225**
- mark-paragraph command 44, **225**
- mark-rectangle command 63, **225**
- mark-rectangle-expands variable 328
- mark_spot primitive 444
- mark_to_column primitive 328, 476
- mark-unhighlights variable 62, 224, 225, 328
- mark-whole-buffer command 59, 61, **225**
- _MATCH_BUF textual macro 570
- Matchdelim buffer variable 86, 329
- matchend primitive 72, 329, 446
- matches_at() subroutine 447
- matches_at_length() subroutine 447
- matches_in() subroutine 447
- matchstart primitive 72, 329, 446
- MAX_CHAR textual macro 521
- max-initial-windows variable 9, 329
- MAXKEYS textual macro 579
- maybe_break_this_line() subroutine 588
- maybe_ding() subroutine 522
- maybe_indent_rigidly() subroutine 476
- maybe_refresh() primitive 471, 578
- maybe-show-matching-delimiter command 44, **225**
- mem_in_use primitive 329, 543
- memcmp() primitive 542
- memcpy() primitive 542
- memfcmp() primitive 542
- memset() primitive 542
- mention() primitive 479, 581
- mention-delay variable 157, 329, 479
- menu bar 32
- menu bar, customizing 33
- menu-bar-flashes variable 33, 329
- menu-bindings variable 33, 329
- menu_command primitive 329, 565
- menu-file variable 33, 329
- menu-stays-after-click variable 33, 330
- menu_tab primitive 570
- menu-width variable 29, 330
- menu-window variable 330
- merge-diff command 55, 56, **225**
- merge-diff-var variable 330
- message-history-size variable 330
- meta characters 474
- Microsoft Visual Studio, integrating with 147
- MicroSpell program 83
- middle_init() subroutine 552
- minimal-coloring variable 114, 330
- minor modes 26
- misc-language-fill-mode variable 330
- MIXEDCASEDRIVES environment variable 126, 127
- mkdir() primitive 508
- mode 25
 - defining a new 584
 - major 25
 - minor 26
- mode line 24, 470
- mode line format 331
- mode_auto_show_delimiter buffer variable 586

mode_default_settings() subroutine 588
 mode-extra buffer variable 330
 mode_extra variable 470
 mode-format variable 24, 107, 116, 331, 470
 mode_keys primitive 580, 581
 mode-line color class 114, 495
 mode-line-at-top variable 332
 mode-line-position variable 332
 mode-more-extra variable 332
 mode_move_level variable 452
 mode_restrict_break buffer variable 589
 MODFOLD textual macro 448
 modified primitive 332, 502
 modified_buffer_region() primitive 458
 modify_region() primitive 483
 monochrome primitive 332, 493
 mouse button, third 31
 mouse support 31
 mouse_auto_off primitive 332, 562
 mouse_auto_on primitive 333, 562
 mouse_buttons() primitive 562
 mouse-center command 175, 176, **225**
 mouse-center-yanks variable 32, 175, 225, 333
 mouse_cursor primitive 563
 MOUSE_DBL_LEFT textual macro 560
 mouse-dbl-selects buffer variable 333, 564
 mouse_display primitive 333, 562
 mouse_handler window variable 563
 MOUSE_LEFT_DN textual macro 560
 mouse_mask primitive 333, 559
 mouse-move command 175, 176, **226**
 mouse-pan command 175, 176, **226**
 mouse_panning primitive 333, 565
 mouse_panning_rate() primitive 565
 mouse_pixel_x primitive 333, 561
 mouse_pixel_y primitive 334, 561
 mouse_pressed() primitive 562
 mouse_screen primitive 334, 463, 560
 mouse-select command 175, **226**
 mouse-selection-copies variable 32, 62, 334
 mouse_shift primitive 334, 561
 mouse-to-tag command 175, 176, **226**
 mouse-update-during-scroll variable 302, 334
 mouse_x primitive 334, 560
 mouse_y primitive 334, 560
 mouse-yank command 175, 176, **226**
 move_by_lines() primitive 451
 move_level() subroutine 452

move_line_to_buffer() subroutine 442
 move_to_column() primitive 475
 move-to-window command 109, 110, **226**
 moving around 42, 106
 moving text 59
 moving windows 31
 MRAUTODEL textual macro 484
 MRCOLOR textual macro 483
 MRCONTROL textual macro 484
 MREND textual macro 484
 MRENDCOL textual macro 484
 MRSTART textual macro 483
 MRSTARTCOL textual macro 484
 MRTYPE textual macro 483
 mshelp2-collection variable 104, 334
 mspellcmd.exe file 83
 muldiv() primitive 480
 multitasking 150
 must_build_mode primitive 334, 471
 must_color_through variable 489
 must_find_func_name variable 587
 MUST_MATCH textual macro 568

N

-n command line flag 398
 name table 543
 name_color_class() primitive 495
 name_debug() primitive 552
 name_help() primitive 554
 name-kbd-macro command 156, 157, 175, **227**
 name_macro() primitive 582
 name_match() primitive 571
 name_name() primitive 544
 name_to_bufnum() primitive 456
 name_type() primitive 544
 name_user() primitive 546
 named pipes, ignoring 48
 named-command command 157, 158, 227
 Narrow, in mode line 176
 narrow_end primitive 335, 444
 narrow_position() subroutine 445
 narrow_start primitive 335, 444
 narrow-to-region command 176, 227, 261
 narrowed_search() subroutine 448
 national characters 136
 national-keys-not-alt variable 335
 near-pause variable 44, 202, 335

need-rebuild-menu variable 335
 NET_DONE textual macro 517, 531, 533
 NET_LOG_DONE textual macro 517
 NET_LOG_WRITE textual macro 517
 net_prompt_get_dir_helper() subroutine 135
 net-prompt-get-directory variable 135, 335
 net-prompt-pattern variable 135, 335, 336
 NET_RECV textual macro 517, 531, 533
 NET_SEND textual macro 517, 531
 new-buffer-translation-type variable 123, 336, 499
 new-c-comments buffer variable 104, 336
 new-file command 117, 118, **227**
 new-file-ext variable 117, 336
 new_file_io_converter variable 501
 new-file-mode variable 117, 336
 new_file_read() primitive 495
 new_file_write() primitive 497, 501
 new-search-delay variable 336
 new_table() primitive 580
 new_variable() primitive 548
 next-buffer command 117, 118, **227**
 next_dialog_item() primitive 577
 next-difference command 56, **227**
 next-error command 153, 155, **227**
 next-match command 49, **228**
 next-page command 107, 176, **228**
 next-position command 49, 153, 155, **228**
 next_screen_line() primitive 472
 next-tag command 52, **228**
 next_user_window() subroutine 460
 next-window command 110, **228**
 nl_forward() primitive 451
 nl_reverse() primitive 451
 NO_MODE_LINE textual macro 470
 no_popup_errors primitive 503
 -nodde command line flag 15
 -noinit command line flag 15
 -nologo command line flag 15
 non-english characters 136
 NONE_OK textual macro 568
 normal-character command 44, 57, 58, 228, 236
 normal-cursor variable 112, 337
 normal-gui-cursor variable 111, 337
 normal_on_modify() subroutine 457
 NOSAVE_ALL textual macro 275
 NOSAVE_AUTOSAVE textual macro 275
 NOSAVE_BUFED textual macro 275

NOSAVE_SAVEALL textual macro 275
 NOSAVE_SESSION textual macro 275
 NOSAVE_WARN textual macro 275
 -noserver command line flag 15, 146
 note() primitive 478
 noteput() primitive 478
 NSIS mode 93
 nsis-auto-show-delim-chars variable 337
 nsis-mode command 93, **229**
 NSS shared files 5
 NT_AUTOLOAD textual macro 550
 NT_AUTOSUBR textual macro 550
 NT_BUFVAR textual macro 544, 548
 NT_BUILTVAR textual macro 544, 545
 NT_COLSCHEME textual macro 492, 544, 548
 NT_COMMAND textual macro 544
 NT_MACRO textual macro 544
 NT_SUBR textual macro 544
 NT_TABLE textual macro 544
 NT_VAR textual macro 544, 548
 NT_WINVAR textual macro 544, 548
 NTFS streams 139
 null, searching for 67
 NUMALT() textual macro 557
 number_of_color_classes() primitive 495
 number_of_popups() primitive 461
 number_of_user_windows() subroutine 461
 number_of_windows() primitive 461
 numbers, entering interactively 159
 NUMCTRL() textual macro 557
 NUMDIGIT() textual macro 557
 NUMDOT textual macro 557
 NUMENTER textual macro 557
 numeric argument 27, 155
 numeric constant 401
 NUMSHIFT() textual macro 557
 numtoi() subroutine 572
 numtoi_with_base() subroutine 573

O

-o command line flag 398
 Objective-C language 88
 octal constant 401
 oem_file_converter() subroutine 500
 oem-to-ansi command 124, 229
 ok_file_match() subroutine 513
 oldkeys.h header file 558, 580

- on, eel keyword 424, 429, 430
- on_exit, eel keyword 415, 416
- on_modify() subroutine 457, 476
- one-window command 109, **229**
- one_window_to_dialog() subroutine 577
- online documentation 38
- only-file-extensions variable 30, 126, 337
- Open With Epsilon shell extension 149
- open-line command 58, **229**
- operator precedence in EEL 417
- opsys primitive 337, 520
- orig_screen_color() primitive 494
- original_argv() primitive 551
- OS_DOS textual macro 337
- OS_OS2 textual macro 337
- OS_UNIX textual macro 337
- OS_WINDOWS textual macro 337
- _our_color_scheme variable 493
- _our_gui_scheme variable 492
- _our_mono_scheme variable 493
- _our_unixconsole_scheme variable 493
- over-mode buffer variable 57, 338
- overwrite-cursor variable 112, 338
- overwrite-gui-cursor variable 111, 338
- overwrite-mode command 57, 58, **229**
- owitheps.dll file 149

P

- p command line flag 15, 146, 398
- page-left command 108, **229**
- page-right command 108, **229**
- page_setup_dialog() primitive 526
- Pager, in mode line 26
- paging 338
- paging-centers-window variable 338
- paging-retains-view variable 338, 465
- paragraphs 43
 - filling 77
- parenthesis matching 44
- parse_string() primitive 447
- parse_url() subroutine 518
- password, typing in a buffer 131
- PASSWORD_PROMPT textual macro 568
- passwords in URLs 136
- PATH environment variable 12
- path, searching for files on a 513
- PATH_ADD_CUR_DIR textual macro 514

- PATH_ADD_EXE_DIR textual macro 514
- PATH_ADD_EXE_PARENT textual macro 514
- path_list_char primitive 338, 512
- PATH_PERMIT_DIRS textual macro 514
- PATH_PERMIT_WILDCARDS textual macro 514
- path_sep primitive 338, 510, 511
- pattern, searching for a 66
- pause-macro command 157, **230**
- PBORDERS textual macro 466
- per-directory file variables 129
- perform_unicode_conversion() primitive 501
- Perl mode 94
- perl-align-contin-lines variable 94, 338
- perl-auto-show-delim-chars variable 339
- perl-brace-offset variable 94, 339
- perl-closeback variable 94, 339
- perl-comment color class 94
- perl-constant color class 94
- perl-contin-offset variable 94, 339
- perl-detect-expression-pattern variable 339
- perl-function color class 94
- perl-heredoc-syntax variable 94, 339
- perl-indent buffer variable 94, 339
- perl-keyword color class 94
- perl-label-indent variable 94, 339
- perl-mode command 94, **230**
- perl-reindent-previous-line variable 340
- perl-string color class 94
- perl-tab-override variable 94, 340
- perl-top-braces variable 94, 340
- perl-top-contin variable 94, 340
- perl-top-struct variable 94, 340
- perl-topindent variable 94, 340
- perl-variable color class 94
- perldoc command 94, 102, 104, **230**
- perldoc-command variable 340
- permanent-menu variable 340
- PERMIT_RESIZE_KEY textual macro 340, 565
- PERMIT_SCROLL_KEY textual macro 340, 565
- PERMIT_WHEEL_KEY textual macro 340, 566
- permit_window_keys primitive 340, 565
- phoneticize_lines() primitive 456
- PHORIZBORDCOLOR textual macro 466
- PHP mode 95
- php-align-contin-lines variable 95, 341
- php-auto-show-delim-chars variable 341
- php-brace-offset variable 95, 341
- php-closeback variable 95, 341

- php-comment-style variable 95, 341
- php-contin-offset variable 95, 341
- php-heredoc-syntax variable 95, 341
- php-indent buffer variable 95, 342
- php-label-indent variable 95, 342
- php-mode command 95, **230**
- php-reindent-previous-line variable 95, 342
- php-tab-override variable 342
- php-top-braces variable 95, 342
- php-top-contin variable 95, 342
- php-top-level-indent buffer variable 95, 342
- php-top-struct variable 95, 342
- php-topindent variable 95, 342
- PIPE_CLEAR_BUF textual macro 532
- PIPE_KEEP_ENV textual macro 533
- PIPE_NOREFRESH textual macro 533
- PIPE_SKIP_SHELL textual macro 533
- PIPE_SYNCH textual macro 532
- pipe_text() subroutine 532
- plink ssh client 133
- pluck-tag command 50, 52, **230**
- point 24
- point primitive 343, 439
- point_spot primitive 444
- pointer to function 543
- pointer to struct, vs. struct 527
- pointer_to_index() primitive 547
- pointers, internal structure 546
- POP_UP_PROMPT textual macro 568
- popup_border color class 495
- popup_near_window() subroutine 468
- popup_title color class 495
- position 440
- position-window-on-screen-line buffer variable 343
- post_compile_hook subroutine 190
- PostScript mode 95
- postscript-auto-show-delim-chars variable 343
- postscript-mode command 95, **230**
- pre_compile_hook subroutine 190
- precedence 417
- prefix keys 157
 - unbinding 166
- prefix-fill-paragraph command 78, **231**
- prepare_url_operation() subroutine 518
- prepare_windows() subroutine 470
- preprocessor lines, moving by 89
- preserve-filename-case variable 127, 343
- preserve-session variable 16, 144, 343
- preserve-session-flags variable 344
- prev_cmd primitive 344, 535, 578
- prev_dialog_item() primitive 577
- prev_forget_buf() subroutine 467
- prev_indenter() subroutine 477
- prev_screen_line() primitive 472
- previous-buffer command 117, 118, **231**
- previous-difference command 56, **231**
- previous-error command 153, 155, **231**
- previous-match command 49, **231**
- previous-page command 107, 176, **231**
- previous-position command 49, 153, 155, **231**
- previous-tag command 52, **232**
- previous-window command 110, **232**
- primary selection in X11 61
- primitive 439
- print-buffer command 137, 138, **232**
- print-buffer-no-prompt command 137, **232**
- print-color-scheme variable 138, 344
- print-destination variable 138, 232, 344
- print-destination-unix variable 138, 232, 344
- print-doublespaced variable 138, 344
- print_eject() primitive 526
- print-heading variable 138, 345
- print-in-color variable 138, 345
- print_line() primitive 526
- print-line-numbers variable 111, 138, 345
- print-long-lines-wrap variable 345
- print-region command 137, 138, **232**
- print-setup command 137, 138, **233**
- print-tabs variable 138, 345
- print_window() primitive 526
- Printf-style format strings 480
- printing 137
- printing variables 159
- PROC_STATUS_RUNNING textual macro 532, 533
- process-backward-kill-word command 152, **233**
- process-coloring-rules variable 151, 345
- process-complete command 152, **233**
- process-completion-dircmds variable 151, 345
- process-completion-style variable 151, 233, 346
- process-completion-windows-programs variable 346
- process-current-directory buffer variable 531
- process-current-directory primitive 346, 508
- process-echo buffer variable 151, 346
- process-enter command **233**

process-enter-whole-line variable 347
 process_exit_status buffer variable 531, 533
 process-exit-status primitive 347
 process_input() primitive 531
 PROCESS_INPUT_CHAR textual macro 531
 PROCESS_INPUT_LINE textual macro 531
 process_kill() primitive 532
 process-max-size buffer variable 347
 process-mode command **233**
 process-next-cmd command 132, 136, 152, **233**
 process-next-error-options variable 347
 process-output-to-window-bottom variable 347
 process-pass-drive-directories variable 347
 process-previous-cmd command 132, 136, 152, **234**
 process-prompt-pattern variable 151, 347
 process_send_text() primitive 532
 process-spell-word-pattern variable 348
 process-tab-size variable 348
 process-view-error-lines variable 153, 348
 process-warn-on-exit variable 348
 process-warn-on-killing variable 348
 process-yank command 152, **234**
 process-yank-confirm variable 152, 349
 profile command 169, **234**
 profiling primitives 554
 programs, running 149
 prompt_box() subroutine 577
 prompt_comp_read() subroutine 569
 prompt-with-buffer-directory variable 126, 349
 prompts, for file names 125
 prox_line_search() subroutine 450
 psftp ssh client 133
 PTEXTCOLOR textual macro 466
 PTITLECOLOR textual macro 466
 ptrlen() primitive 546
 pull-highlight color class 114
 pull-word command 102, 114, **234**
 pull-word-from-tags variable 101, 349
 pull-word-fwd command 102, **234**
 push command 149, 150, **235**
 push-cmd variable 154, 349
 push-cmd-unix-interactive variable 350
 push-cmd-unix-macos-interactive variable 350
 put_directory() subroutine 508
 putenv() primitive 520
 PuTTY ssh client 133
 PVERTBORDCOLOR textual macro 466
 Python mode 95

python-auto-show-delim-chars variable 350
 python-comment-region-options variable 285, 350
 python-delete-hacking-tabs variable 96, 350
 python-indent buffer variable 96, 350
 python-indent-to-comment variable 350
 python-indent-with-tabs variable 96, 350
 python-indenter command **234**
 python-language-level variable 96, 350
 python-mode command 96, **235**
 python-tab-override variable 96, 351

Q

-q command line flag 398
 QUERY textual macro 448
 query-replace command 64, 65, **235**
 quick_abort() primitive 534, 581
 quick-dired-command command 140, **235**
 -quickup command line flag 16
 quiet-write-state variable 351
 quit_bufed() subroutine 467
 quoted-insert command 58, **236**
 quoting special chars in searches 45

R

-r command line flag 16, 552
 R mode 96
 r-mode command 96, **236**
 raw_xfer() primitive 441
 re_compile() primitive 447
 RE_FIRST_END textual macro 447
 RE_FORWARD textual macro 447
 RE_IGNORE_COLOR textual macro 447
 re_match() primitive 447
 RE_REVERSE textual macro 447
 re_search() primitive 446, 447
 RE_SHORTEST textual macro 447
 _read_aborted variable 496
 read_file() subroutine 496
 read-only files 504
 read-only files and buffers 119
 read-session command 145, 146, **236**
 readme.txt file 19
 readonly-pages variable 26, 119, 351, 458
 readonly-warning variable 120, 351
 realloc() primitive 542
 rebuild-menu command 33, **237**
 recall-id variable 351

- recall-longest-response variable 351
- recall-maximum-session variable 351
- recall-maximum-size variable 352
- recall-prior-response-options variable 352
- recall-response-selected variable 352
- recalling previous commands 30
- recognize-password-pattern variable 131, 352
- recognize-password-prompt buffer variable 131, 352
- recolor_by_lines() subroutine 490
- recolor_from_here variable 490
- recolor_from_top() subroutine 490
- recolor_partial_code() subroutine 492
- recolor_range variable 489
- record-customizations variable 168, 353, 387
- record-kbd-macro command **237**
- recording-suspended variable 353
- rectangle editing 63
- rectangle_standardize() primitive 486
- _recursion_level variable 536
- recursive_edit() subroutine 536
- recursive_edit_preserve() subroutine 536
- redisplay command 156, 157, **237**
- redo command 105, 106, **237**
- redo vs. redo-changes 105
- redo-by-commands command 106, **237**
- redo-changes command 105, 106, **237**
- redo-movements command 106, **237**
- refresh() primitive 471
- refresh-files command **237**
- reg_tab primitive 580
- REGEX textual macro 448
- regex-first-end variable 72, 353
- regex-replace command 65, 74, 75, **238**
- regex_replace_in_string() subroutine 540
- regex-search command 48, 74, 75, **238**
- regex-shortest variable 72, 353
- REGINCL textual macro 483
- region 59
- region_type() subroutine 484
- REGLINE textual macro 483
- REGNORM textual macro 483
- REGRECT textual macro 483
- regular expression assertions 73
- regular expressions 45, 66, 446
- reindent-after-c-yank variable 86, 353
- reindent-after-perl-yank variable 94, 353
- reindent-after-vbasic-yank variable 100, 354
- reindent-after-yank variable 79, 262, 354
- reindent-c-comments buffer variable 354
- reindent-c-preprocessor-lines buffer variable 87, 354
- reindent-one-line-c-comments buffer variable 354
- reindent-perl-comments buffer variable 94, 354
- reject_client_connections primitive 523
- relative() primitive 510
- release-notes command 38, 39, **239**
- remote_dirname_absolute() subroutine 509
- remote_file_type() subroutine 512
- remove_final_view() subroutine 465
- remove_line_highlight() subroutine 486
- remove_region() primitive 483
- remove_window() primitive 459
- rename-buffer command **239**
- rename_file() primitive 503
- renaming commands or variables 162
- renaming files 141
- repeating commands 27
- repeating, keys 556
- Repeating: Numeric Arguments 27
- replace() primitive 440
- replace-again command **239**
- replace-by-case variable 65, 355
- REPLACE_FUNC() textual macro 545
- replace_in_existing_hook() subroutine 450
- replace_in_readonly_hook() subroutine 451
- replace-in-region variable 355
- replace_name() primitive 544, 545
- replace-num-changed variable 355, 448
- replace-num-found variable 355, 448
- replace-string command 64, 65, **239**
- replacing in multiple files 65
- reserved EEL keywords 401
- reset-mode command **239**
- reset_modified_buffer_region() primitive 458
- resize-menu-list variable 355
- resize-rectangle-on-tab variable 355
- resize_screen() primitive 473
- resizing windows 31
- restart-concurrent variable 154, 355
- restore-color-on-exit variable 114, 356, 494
- restore_screen() subroutine 462
- restore_vars() primitive 415, 416
- resume-client command 147, **239**
- resynch-match-chars variable 54, 356
- retag-files command 52, **240**

- return, eel keyword 414, 415
- return-raw-buttons variable 356, 577
- rev-search-key variable 47, 356
- REVERSE textual macro 448
- reverse-incremental-search command 48, **240**
- reverse-regex-search command 48, 75, **240**
- reverse-replace command 65, **240**
- reverse-search-again command 47, 48, **240**
- reverse-sort-buffer command 75, 76, **240**
- reverse-sort-region command 75, 76, **240**
- reverse_split_string() subroutine 518
- reverse-string-search command 46, 48, **240**
- revert-file command 119, **240**
- reverting to old file 119
- rgb_to_attr() primitive 494
- right margin wrap 77
- right_align_columns() subroutine 454
- rindex() primitive 539
- rmdir() primitive 508
- rmode-auto-show-delim-chars variable 356
- rmode-indent buffer variable 356
- rmode-indent-to-comment variable 356
- rmode-tab-override variable 356
- RO, in mode line 119
- root_keys primitive 580, 581
- ROWARN_BELL textual macro 351
- ROWARN_BUF_RO textual macro 351
- ROWARN_GREP textual macro 351
- ROWARN_MSG textual macro 351
- ROWARN_REFRESH textual macro 351
- ruler, displaying 111
- run-by-mouse variable 356, 564
- run-ssh-agent.bat file 133
- run_topkey() subroutine 581
- run_viewer() primitive 534
- run-with-argument command 155, **240**
- running other programs 149

S

- s command line flag 16, 398
- safe_copy_buffer_variables() subroutine 548
- save-all-buffers command 121, 154, **241**
- save-all-without-asking variable 357
- save-file command 121, **241**, 261
- save_remote_file() subroutine 500
- save_screen() subroutine 462
- save_spot, eel keyword 415, 416, 443

- save_state() primitive 550
- save_var, eel keyword 414, 415, 416
- save-when-making variable 154, 357
- save_without_prompt variable 502
- saving customizations 162
- saving files automatically 121
- say() primitive 477, 479, 481, 534, 541
- sayput() primitive 477, 479, 481
- SCON_COMPARE textual macro 450
- SCON_RECORD textual macro 449
- SCON_RESTORE textual macro 449
- scope of variables 402
- scp-client-style variable 133, 135, 357
- scp-list-flags variable 133, 357
- scp-read-file-no-user-template variable 357
- scp-read-file-template variable 358
- scp-run-helper-no-user-template variable 358
- scp-run-helper-template variable 358
- scp-unix-sftp-command variable 358
- scp-windows-sftp-command variable 133, 359
- scp-write-file-no-user-template variable 359
- scp-write-file-template variable 359
- scratch buffers 60
- screen 23
- screen-border color class 114
- screen_cols primitive 359, 472
- screen-decoration color class 114
- screen_lines primitive 359, 472
- screen_messed() primitive 472
- screen_to_window() primitive 464
- screen_to_window_id() primitive 523
- scripts
 - complex 137
- scroll bar 31
- Scroll Lock key 106
- scroll-at-end variable 42, 360
- scroll_bar_line() primitive 564
- scroll-bar-type variable 360
- scroll_by_wheel() subroutine 564
- scroll-down command 107, **241**
- scroll-init-delay variable 31, 360
- scroll-left command 108, **241**
- scroll-rate variable 31, 360
- scroll-right command 108, **241**
- scroll-up command 107, **241**
- scrollbar_handler() subroutine 565
- scrolling, lines 108
- search() primitive 446

- search-again command 47, 48, **241**
- search_continuation primitive 449
- search-defaults-from variable 47, 65, 360
- search-delete-match variable 360
- search-in-menu variable 29, 360
- search-in-region variable 242, 361
- search-man-pages command 102, **241**
- search-man-pages-shows-all variable 361
- search-positions-at-start variable 361
- search_read() subroutine 449
- search-region command 47, 48, **242**
- search-wraps variable 46, 361
- searching
 - and replacing 64
 - case folding 45
 - conventional 46
 - for special characters 45
 - for words 45
 - incremental 45
 - incremental mode 46
 - regular expression 45
- searching multiple files 48
- secure shell 132
- see-delay variable 111, 361, 478
- select-and-grep command 49, **242**
- select-browse-file command 53, 54, **242**, 274
- select-buffer command 117, 141, **242**
- select_low_window() primitive 462
- select_menu_item() subroutine 573
- select_printer() primitive 526
- select-tag-file command 51, 52, **243**
- selectable-colors variable 361
- selected_color_scheme primitive 275, 361, 380, 492
- Send To menu, putting Epsilon on a 148
- send-invisible command **243**
- sendeps program 148
- sendnew.exe file 148
- sendonly command line flag 16
- sentence commands 43
- sentence-end variable 361
- sentence-end-double-space variable 43, 361
- server command line flag 16, 146
- server-raises-window variable 147, 362, 523
- session-always-restore variable 145, 362
- session-default-directory variable 145, 362
- session-file-name variable 146, 362
- session-restore-biggest-file variable 362
- session-restore-biggest-remote-file variable 362
- session-restore-directory variable 145, 362
- session-restore-directory-buffers variable 363
- session-restore-exclude-file-types variable 363
- session-restore-files variable 145, 363
- session-restore-max-directories variable 363
- session-restore-max-files variable 144, 363
- session-tree-root variable 145, 363
- session-warn-when-saving variable 364
- sessions, restoring 144
- set-abort-key command 106, **243**
- set-any-variable command 159, 162, **243**
- set-bookmark command 50, 51, 53, **243**
- set-bookmark-automatically variable 50, 243, 364
- set_buf_modified() subroutine 502
- set_buf_point() subroutine 457
- set_buffer_filename() subroutine 502
- set_character_color() primitive 441, 487
- set_character_property() primitive 538
- set_chars() primitive 541
- set-color command 114, 115, **243**, 257
- set_color_pair() primitive 493
- set-comment-column command 105, **244**
- set-debug command 169, **244**
- set-dialog-font command 113, **244**
- set-display-characters command 111, 112, 244, 474
- set_display_func_name() subroutine 587
- set-display-look command 116, **245**
- set-encoding command 137, **245**
- SET_ENCODING() textual macro 499
- set-file-name command 121, **245**
- set_file_opsys_attribute() primitive 504
- set_file_read_only() primitive 504
- set-fill-column command 78, **245**
- set-font command 112, 113, **246**
- set-line-translate command 123, 124, **246**
- set_list_keys() subroutine 580
- set-mark command 60, **246**
- set_mode() subroutine 470
- set_mode_message() subroutine 330, 470
- set_name_debug() primitive 552
- set_name_help() primitive 554
- set_name_user() primitive 546
- set-named-bookmark command 50, **246**
- set_num_var() primitive 545
- set-printer-font command 113, 137, **246**

- set_range() primitive 579
- set_region_type() subroutine 484
- set-show-graphic command 111, 112, **246**
- set_spot() subroutine 444
- set_str_var() primitive 545
- set_swapname() primitive 543
- set-tab-size command 110, 112, **247**
- set_tagged_region() primitive 488
- SET_TRANSLATE() textual macro 499
- set-variable command 45, 159, 162, **247**, 545, 546
- set-variable, in command file 166
- set-want-backup-file command **247**
- set_wattrib() primitive 466
- set_window_caption() primitive 578
- setjmp() primitive 536
- setting
 - colors 113
 - variables 159
- setting bookmarks 50
- sftp program 132
- shebang line 127
- SHELL environment variable 12, 150
- shell extension, Open with Epsilon 149
- shell mode 96
- shell() primitive 529, 530
- shell-auto-show-delim-chars variable 364
- shell-command command 150, **247**
- shell-heredoc-syntax variable 96, 364
- SHELL_HIDE textual macro 530
- SHELL_KEEP_ENV textual macro 530, 534
- SHELL_MAXIMIZED textual macro 530
- SHELL_MINIMIZED textual macro 530
- shell-mode command 96, **248**
- SHELL_NO_CD_MAGIC textual macro 531
- SHELL_NO_SYNCH textual macro 530
- SHELL_SYNCH textual macro 530, 534
- shell-tab-override variable 96, 364
- shelling commands 149
- shift_pressed() primitive 562
- shift-selecting variable 364
- shift-selects variable 59, 364
- short, ccl keyword 403, 543
- shortcut, running Epsilon from a 148
- show-all-variables variable 364
- show_binding() subroutine 555
- show-bindings command 37, 38, **248**
- show_char() primitive 558
- show-connections command 130, 132, **248**
- show-last-keys command 38, **248**
- show-matching-delimiter command 44, 86, 89, **248**, 329
- show-menu command 32, **248**
- show_minor_mode_ subroutines 470
- show-mouse-choices variable 365, 564
- show-point command 108, **248**
- show_replace() subroutine 448
- show-spaces buffer variable 111, 188, 365
- show-standard-bitmaps command **248**, 525
- show_status primitive 365, 453
- show-tag-line variable 51, 365
- show_text() primitive 478
- show_url() subroutine 515
- show-variable command 161, 162, **249**
- show-version command **249**
- show-view-bitmaps command **249**, 525
- show-when-idle variable 116, 365, 555
- show-when-idle-column variable 116, 366
- show_window_caption() subroutine 578
- shrink-window command 110, **249**
- shrink-window-horizontally command 110, **249**
- shrink-window-interactively command 110, **249**
- signal_suspend() primitive 522
- simple debugger 391
- simple_re_replace() subroutine 449
- size() primitive 439
- sizeof operator 419
- soft-tab-size buffer variable 79, 181, 211, 212, 366, 369
- sort_another() subroutine 453
- sort-buffer command 75, **249**
- sort-case-fold buffer variable 75, 366
- sort-region command 75, **249**
- sort-tags command 52, **250**
- sorting 75, 453
- source code browsing 62
- Sp, in mode line 81
- SPACE_VALID textual macro 569
- SPACEBAR textual macro 557
- special files, ignoring 48
- spell checking 81
- spell-buffer-or-region command 82, 83, **250**
- spell-configure command 81, 83, **250**
- spell-correct command 82, 83, **250**
- spell-grep command 82, 83, **251**
- spell-helper-program variable 250, 367
- spell_language_prefix primitive 83
- spell-mode command 82, 83, **251**

- split_string() subroutine 518
- split-window command 109, **251**
- split-window-vertically command 109, **251**
- spot 442
- spot, eel keyword 403
- spot_to_buffer() primitive 443
- sprintf() primitive 541
- ssh agents 132, 133
- ssh command 132, 136, **251**
- ssh program 132
- ssh-command-windows variable 133, 367, 368
- ssh-interpret-output variable 132, 367
- ssh-mode command 136, **251**
- ssh-no-user-template variable 133, 368
- ssh-template variable 133, 368
- standard-color variable 493
- standard-gui variable 493
- standard-mono variable 493
- standard_tab_cmd() subroutine 476
- standard-toolbar command 174, **251**, 525
- standardize_remote_pathname() subroutine 509
- start-kbd-macro command 156, 227, **252**
- start-make-in-buffer-directory variable 368
- start_of_function variable 587
- start_print_job() primitive 526
- start-process command 150, 152, **252**, 530
- start-process-in-buffer-directory variable 368
- start_profiling() primitive 554
- start_up() subroutine 545, 552
- starting Epsilon 9
- STARTMATCH textual macro 568
- starts_with_in_list() subroutine 540
- startup files 162
- state file 16, 162
- state_extension primitive 368, 550
- state_file primitive 552
- state-file-backup-name variable 163, 368
- _std_disp_class variable 474
- std_find_it() subroutine 496
- stop-process command 152, **252**, 532
- stop_profiling() primitive 554
- strcat() primitive 541
- strchr() primitive 539
- strcmp() primitive 539
- strcpy() primitive 541
- strfcmp() primitive 539
- stricmp() primitive 539
- string constant 402

- string_matches_pattern() subroutine 540
- string_matches_regex() subroutine 540
- string_replace() subroutine 355, 448, 540
- string-search command 46, 48, **252**
- strings in when_loading() ftns 549
- strkeep() subroutine 542
- strlen() primitive 539
- strncat() primitive 541
- strncmp() primitive 539
- strncpy() primitive 541
- strnfcmp() primitive 539
- strpbrk() subroutine 540
- strpbrk_cnt() subroutine 540
- strrchr() primitive 539
- strsave() primitive 542, 549
- strstr() primitive 539
- strtoi() subroutine 573
- structure-or-union specifier 407
- stuff() primitive 440
- stuff_macro() subroutine 558
- subroutine 404
- suffix_ subroutines 84, 585
- suffix_default() subroutine 585
- suffix_none() subroutine 585
- Susp, in mode line 156
- suspend-epsilon command 144, **252**
- swap file 14
- switch, eel keyword 413, 415
- switch-buffers command 118, **252**, 369
- switch-buffers-options variable 252, 369
- switch_to_buffer() subroutine 467
- switch-windows command **253**
- switches
 - for EEL 397
 - for Epsilon 13
- symbolic links, ignoring 306
- syntax highlighting 114
- system variables 159
- system_window primitive 369, 466

T

- tab size, setting 110, 321, 369
- tab_convert() subroutine 476
- tab-size buffer variable 79, 87, 110, 181, 211, 212, 366, 369, 473
- tab-width file variable 128
- tabify-buffer command 80, **253**

- tabify-region command 80, **253**
- table_count primitive 369, 581
- table_keys primitive 581
- table_prompt() subroutine 581
- tabs, used for indenting 79
- tag, struct or union 408
- tag-ask-before-retagging variable 51, 369
- tag-batch-mode variable 369
- tag-by-text variable 52, 369
- tag-c-preprocessor-skip-pat variable 369
- tag-case-sensitive variable 51, 370
- tag-declarations variable 51, 370
- tag-display-width variable 370
- tag-files command 51, 52, **253**
- tag-list-exact-only variable 370
- tag-modifier-kwd variable 370
- tag-options variable 370
- tag-pattern-c variable 370
- tag-pattern-default variable 370
- tag-pattern-perl variable 370
- tag-relative variable 52, 371
- tag-show-percent variable 371
- tag_suffix_default() subroutine 519
- tag_suffix_none() subroutine 519
- tag-which-items variable 371
- tagged regions 488
- tagging function names 50
- TB_BORD() textual macro 463
- Tcl mode 97
- tcl-auto-show-delim-chars variable 371
- tcl-indent buffer variable 97, 371
- tcl-indent-and-show-matching-delimiter command **253**
- tcl-indent-cmd command **253**
- tcl-mode command 97, **253**
- teach command line flag 16
- telnet command 131, 132, **253**
- Telnet URL 131
- telnet_host() primitive 514
- telnet_id variable 515
- telnet-interpret-output variable 131, 371
- telnet-mode command 132, **254**
- telnet_send() primitive 514
- telnet_server_echoes() primitive 515
- TEMP environment variable 12, 14
- temp_buf() subroutine 457
- template, file name 122, 513
- term_clear() primitive 482
- term_cmd_line() subroutine 473
- term_init() subroutine 473
- term_mode() subroutine 473
- term_position() primitive 482
- term_write() primitive 482
- term_write_attr() primitive 482
- terminal program under X11 17
- terminal-epsilon program 5
- TeX mode 97
- tex-auto-fill-mode variable 372
- tex-auto-show-delim-chars variable 372
- tex-boldface command 99, **254**
- tex-center-line command 99, **254**
- tex-close-environment command 99, **254**
- tex-display-math command 99, **254**
- tex-environment command 99, **254**
- tex-environment-name variable 372
- tex-fill-paragraph command **254**
- tex-footnote command 99, **254**
- tex-force-latex buffer variable 97, 372
- tex-force-quote command 99, **254**
- tex-inline-math command 99, **255**
- tex-italic command 99, **255**
- tex-left-brace command 99, **255**
- tex-look-back variable 97, 372
- tex-math-escape command 99, **255**
- tex-mode command 100, **255**
- tex-mouse-select command **255**
- tex-paragraphs buffer variable 43, 204, 372
- tex-quote command 99, **255**
- tex-rm-correction command 99, **255**
- tex-save-new-environments variable 372
- tex-slant command 99, **256**
- tex-small-caps command 99, **256**
- tex-spell-options variable 372
- tex-typewriter command 99, **256**
- text color class 114, 495
- text_color primitive 372, 495
- text_height() primitive 461
- text_width() primitive 461
- third mouse button 31
- this-buffer-grep command 49, **256**
- this_cmd primitive 373, 535, 578
- tilde-border variable 373
- tilde_only() subroutine 467
- tilde-scroll-bar variable 373
- time_and_day() primitive 527
- time_begin() primitive 527
- time_done() primitive 527

`time_ms()` primitive 527
`time_remaining()` primitive 527
 TIMER, type definition 527
 title, of window 469
 TITLECENTER textual macro 469
 TITLELEFT() textual macro 469
 TITLERIGHT() textual macro 469
 TMP environment variable 14
`tmp_buf()` subroutine 457
`to_another_buffer()` subroutine 467
`to_begin_line()` textual macro 451
`to_buffer()` subroutine 466
`to_buffer_num()` subroutine 466
`to_column()` subroutine 475
`to_end_line()` textual macro 451
 to-indentation command 80, **256**
 to-left-edge command **256**
 to-right-edge command **256**
`to_virtual_column()` subroutine 476
 toggle-borders command 115, 256
 toggle-case-fold command 45, 48, **257**
 toggle-menu-bar command 32, **257**
 toggle-scroll-bar command 31, **257**
 toggle-toolbar command 174, 175, **257**
`tokenize_lines()` primitive 455
`tolower()` primitive 537
`toolbar_add_button()` primitive 525
`toolbar_add_separator()` primitive 525
`toolbar_create()` primitive 525
`toolbar_destroy()` primitive 525
`top_level` primitive 537
 Topindent buffer variable 86, 373
 TOPLEFT textual macro 461
`toupper()` primitive 537
 tracing debugger 169
 translation 122, 498
 translation-type buffer variable 123, 373, 495, 497, 499, 500
 transpose-characters command 77, **257**
 transpose-lines command 77, **257**
 transpose-words command 77, **257**
 transposing things 76
`try_calling()` primitive 544
`try_show_url()` subroutine 515
`#tryinclude` preprocessor command 399
 tutorial 9
 tutorial command **257**
`two_scroll_box()` subroutine 577

type names 409
 type point 152
 type specifier 404
 TYPE_BYTE textual macro 547, 548
 TYPE_CARRAY textual macro 547, 548
 TYPE_CHAR textual macro 547, 548
 TYPE_CPTR textual macro 547, 548
 TYPE_INT textual macro 547, 548
 TYPE_INT32 textual macro 547, 548
 TYPE_OTHER textual macro 547
 type_point primitive 373, 531
 TYPE_POINTER textual macro 547
 TYPE_SHORT textual macro 547, 548
 typing-deletes-highlight variable 60, 360, 373
 typing-hides-highlight variable 60, 373

U

unbind-key command 157, 158, **258**
 UNC network file syntax 139
 uncompress-files variable 119, 374
`#undef` preprocessor command 399
 underlined text 112
 undo command 105, 106, 237, **258**
 undo vs. undo-changes 105
 undo-by-commands command 106, **258**
 undo-changes command 105, 106, **258**
`undo_count()` primitive 445
 UNDO_DELETE textual macro 445
 UNDO_END textual macro 445
`undo_flag` primitive 374, 445
 UNDO_FLAG textual macro 445
 UNDO_INSERT textual macro 445
 undo-keeps-narrowing buffer variable 374
 UNDO_MAINLOOP textual macro 445
`undo_mainloop()` primitive 445, 578
 UNDO_MOVE textual macro 445
 undo-movements command 106, **258**
`undo_op()` primitive 445
 UNDO_REDISP textual macro 445
`undo_redisplay()` primitive 445
 UNDO_REPLACE textual macro 445
 undo-size buffer variable 106, 374
`ungot_key` primitive 374, 555, 558
 Unicode conversion 136, 501
 UNICODE textual macro 398
`unicode_character_name()` primitive 453, 558
 unicode-convert-from-encoding command 124, 137, **258**

- unicode-convert-to-ascii command **258**
- unicode-convert-to-encoding command 124, 137, **258**
- uniform resource locator (URL) 130
- uniq command 76, **259**
- unique_file_ids_match() subroutine 506
- unique_filename_identifier() primitive 506
- Unix files 122
- Unix, Epsilon for 5
- Unix, in mode line 123
- unsaved_buffers() subroutine 502
- unseen_msgs() primitive 478
- unseen_msgs_time() primitive 478
- untabify-buffer command 63, 80, **259**
- untabify-region command 80, **259**
- untag-files command 52, **259**
- up-line command 42, **259**
- update Epsilon 170
- update_readonly_warning() subroutine 496
- updating Epsilon 170
- uppercase-word command 63, 64, **259**
- URL (uniform resource locator) 130
- URL syntax 136
- url_is_a_directory() subroutine 504
- url_operation() subroutine 516
- url_services primitive 512
- use_alternate_dialog variable 577
- use_alternate_dialog_name variable 577
- use_alternate_dialog_tag variable 577
- use-c-macro-rules variable 88, 374
- use_common_file_dialog() subroutine 574
- use_common_file_dlg() subroutine 574
- use-compile-command-file-variable variable 374
- use_default primitive 375, 548
- USE_DEFAULT_COLORS environment variable 113
- use-file-variables variable 129, 375
- use-grep-ignore-file-variables variable 49, 242, 375
- use-grep-ignore-file-variables-all variable 242
- use-process-current-directory variable 375, 508
- user, eel keyword 159, 424, 546
- user_abort primitive 375, 534
- Users directory 13
- using_new_font primitive 482
- using_oem_font() primitive 482
- UTF-16 encoding 136

V

- v command line flag 398

- variables
 - buffer-specific 161, 403, 424
 - in EEL 402
 - setting & showing 159
 - window-specific 161, 403, 424
- varptr() primitive 547
- vartype() primitive 547
- vartype_class() subroutine 548
- VBasic mode 100
- vbasic-auto-show-delim-chars variable 375
- vbasic-indent buffer variable 376
- vbasic-indent variable 100
- vbasic-indent-case variable 376
- vbasic-indent-subroutines variable 100, 376
- vbasic-indent-with-tabs variable 100, 376
- vbasic-language-level variable 100, 376
- vbasic-mode command 100, **259**
- vbasic-reindent-previous-line variable 297, 376
- vc command line flag 16, 472
- vcolor command line flag 16
- verenv() primitive 520
- version primitive 375, 552
- versioned-file-string variable 376
- vert-border color class 114, 495
- VERTICAL textual macro 460
- vf command line flag 16
- VHDL mode 100
- vhdl-auto-show-delim-chars variable 376
- vhdl-indent buffer variable 377
- vhdl-indenter command **259**
- vhdl-mode command 100, **259**
- Vi/Vim file variables 130
- _view_border variable 465
- _view_bottom variable 465
- view_buf() subroutine 464, 570
- view_buffer() subroutine 464
- _view_left variable 465
- view_linked_buf() subroutine 464, 554
- view_loop() subroutine 465
- view-lugaru-web-site command 131, 132, **260**
- view-process command 153, 155, **260**
- _view_right variable 465
- view_tab primitive 570
- _view_title variable 465
- _view_top variable 465
- view-web-site command 131, 132, **260**
- virtual_column() subroutine 476
- virtual-insert-cursor variable 112, 377

virtual-insert-gui-cursor variable 111, 377
 virtual_mark_column() subroutine 476
 virtual-overwrite-cursor variable 112, 377
 virtual-overwrite-gui-cursor variable 111, 377
 virtual-space buffer variable 42, 377
 VisEpsil.dll file 147
 visit-file command 119, **260**
 Visual Basic mode 100
 Visual Studio, integration with 147
 visual-diff command 55, 56, **260**
 visual-diff-mode command 56, **260**
 -vl command line flag 17, 472
 -vmono command line flag 17
 VMS 310
 volatile, eel keyword 411
 -vt command line flag 17
 -vv command line flag 17
 -vx command line flag 17
 -vy command line flag 17

W

-w command line flag 17, 300
 w-bottom variable 377
 w-left variable 377
 w-right variable 377
 w-top variable 377
 -wait command line flag 17
 wait_for_key() primitive 555, 556, 578, 579
 wait_for_unseen_msgs() subroutine 478
 wall-chart command 38, 39, **260**
 want-auto-save variable 122, 268, 378
 want-backups buffer variable 121, 271, 378
 want-bell variable 116, 378, 522
 want-code-coloring buffer variable 114, 378
 want_cols primitive 378, 472
 want-common-file-dialog variable 125, 378
 want-debug-coloring buffer variable 378, 491
 want-display-host-name variable 378
 want-gui-menu variable 378
 want-gui-printing variable 138, 379
 want-gui-prompts variable 29, 379
 want_lines primitive 379, 472
 WANT_MODE_LINE textual macro 470
 want-sorted-tags variable 52, 379
 want-state-file-backups variable 163, 379
 want_toolbar primitive 379, 525
 want-warn buffer variable 121, 379

want-window-borders variable 379
 warn-before-overwrite variable 120, 380
 warn_existing_file() subroutine 498
 was_key_shifted() subroutine 562
 was-quoted variable 380
 Web URL 131
 web-search command **261**
 web-search-url variable 261, 380
 wget 131
 what-is command 38, **261**
 wheel mouse button 31
 wheel mouse support 380
 wheel-click-lines variable 380
 when_aborting() subroutine 534
 when_activity buffer variable 531, 533
 when_displaying variable 491
 when_exiting() subroutine 535
 when_idle() subroutine 555
 when_loading() subroutine 429, 549
 when_net_activity buffer variable 517
 when_repeating() subroutine 556
 when_restoring() subroutine 552
 when_setting_ subroutines 545
 while, eel keyword 413
 widen-buffer command 176, **261**
 wildcard file patterns 138
 wildcard searching 66
 win-askpass.exe file 18
 WIN_BUTTON textual macro 566, 578
 win_display_menu() primitive 524
 WIN_DRAG_DROP textual macro 146, 565
 WIN_EXIT textual macro 565
 win_help_contents() primitive 524
 WIN_HELP_REQUEST textual macro 565
 win_help_string() primitive 524
 win_load_menu() primitive 524
 win_menu_popup() primitive 525
 WIN_MENU_SELECT textual macro 329, 565
 WIN_RESIZE textual macro 565
 WIN_VERT_SCROLL textual macro 565
 WIN_WHEEL_KEY textual macro 340, 564, 566
 window
 keyword 424, 430
 window handle 460
 window number 460
 window storage class 162
 window title 469
 window, eel keyword 162, 403, 424

`window_at_coords()` primitive 463
`window-black` color class 115
`window-blue` color class 115
`window_bufnum` primitive 380, 467
`window-caption` variable 115, 380
`window-caption-buffer` variable 115, 380
`window-caption-file` variable 115, 380
`window_content_width()` primitive 461
`window_create()` subroutine 462
`window_edge()` primitive 461
`window_end` primitive 381, 467
`window_extra_lines()` primitive 468
`_window_flags` window variable 470
`window_handle` primitive 381, 460
`window_height` primitive 381, 461
`window_kill()` primitive 459
`window_left` primitive 381, 464
`window_line_to_position()` primitive 468
`window_lines_visible()` primitive 574
`window_number` primitive 381, 460
`window_one()` primitive 459
`window-overlap` variable 107, 381
`window_scroll()` primitive 468
`window-specific` variables 161, 403
`window_split()` primitive 460
`window_start` primitive 381, 467
`window_title()` primitive 469
`window_to_fit()` subroutine 468
`window_to_screen()` primitive 463
`window_top` primitive 381, 464
`window_width` primitive 381, 461
windows 23

- creating 109
- deleting 109
- selecting 109
- sizing 110

`windows_foreground()` primitive 523
`windows_help_from()` subroutine 524
`windows_maximize()` primitive 523
`windows_minimize()` primitive 523
`windows_restore()` primitive 523
`windows_set_font()` primitive 482
`windows_state()` primitive 523
`winexec()` primitive 533
`winhelp-display-contents` variable 381
`winpty.exe` file 367
`WINSTATE_MAXIMIZED` textual macro 523
`WINSTATE_MINIMIZED` textual macro 523

word commands 42
word searching 45
`WORD` textual macro 448
word wrap mode 77
`word_in_list()` subroutine 540
`word-pattern` buffer variable 43, 204, 382
`word_search()` subroutine 448
`wrap-dired-live-link` variable 143, 382
`wrap-grep` variable 382
`wrap-info-mode` variable 40, 382
`wrap-split-vertically` variable 109, 382
`wrap-view-error-lines` variable 348, 382
wrapping during searches 46
wrapping text as you type 77
wrapping, lines 108
write-file command 121, **261**
write-files-and-exit command **261**
write-line-translate file variable 128
`write_part()` subroutine 503
write-region command 121, **261**
write-session command 145, 146, **261**
write-state command 163, 164, **261**, 387, 424
WWW URL 131

X

`x_pixels_per_char()` primitive 561
X11 windowing system 5
xdvi dvi viewer program 98
`xfer()` subroutine 441
`xfer_rectangle()` subroutine 485
XML mode 91
`xml-asp-coloring` variable 92, 383
`xml-auto-fill-combine` variable 383
`xml-auto-fill-mode` variable 383
`xml-auto-indent` variable 91, 383
`xml-indent` buffer variable 91, 383
`xml-mode` command 93, **262**
`xml-paragraph-elements` variable 383
`xml-php-coloring` variable 92, 384
`xml-reindent-previous-line` variable 384
`xml-sort-by-attribute-name` command 93, **262**
`xml-spell-options` variable 384
xterm 17
`xterm-color` variable 493

Y

`y_pixels_per_char()` primitive 561

yank command 60, 61, 152, **262**
yank-line-retains-position variable 384
yank-options variable 384
yank-pop command 61, **262**
yank-rectangle-to-corner variable 384
yank-x-selection command 62, **262**
yap dvi viewer program 98

Z

zap() primitive 456
zeroed, eel keyword 424
zoom-window command 109, **262**