

# XEmacs Internals Manual

---

Version 1.4, March 2001

**Ben Wing**  
**Martin Buchholz**  
**Hrvoje Niksic**  
**Matthias Neubauer**  
**Olivier Galibert**

---

Copyright © 1992 - 1996, 2001 Ben Wing.  
Copyright © 1996, 1997 Sun Microsystems, Inc.  
Copyright © 1994 - 1998 Free Software Foundation.  
Copyright © 1994, 1995 Board of Trustees, University of Illinois.

Version 1.4  
March 2001.

Permission is granted to make and distribute verbatim copies of this manual provided the copyright notice and this permission notice are preserved on all copies.

Permission is granted to copy and distribute modified versions of this manual under the conditions for verbatim copying, provided also that the section entitled “GNU General Public License” is included exactly as in the original, and provided that the entire resulting derived work is distributed under the terms of a permission notice identical to this one.

Permission is granted to copy and distribute translations of this manual into another language, under the above conditions for modified versions, except that the section entitled “GNU General Public License” may be included in a translation approved by the Free Software Foundation instead of in the original English.

# 1 A History of Emacs

XEmacs is a powerful, customizable text editor and development environment. It began as Lucid Emacs, which was in turn derived from GNU Emacs, a program written by Richard Stallman of the Free Software Foundation. GNU Emacs dates back to the 1970's, and was modelled after a package called "Emacs", written in 1976, that was a set of macros on top of TECO, an old, old text editor written at MIT on the DEC PDP 10 under one of the earliest time-sharing operating systems, ITS (Incompatible Timesharing System). (ITS dates back well before Unix.) ITS, TECO, and Emacs were products of a group of people at MIT who called themselves "hackers", who shared an idealistic belief system about the free exchange of information and were fanatical in their devotion to and time spent with computers. (The hacker subculture dates back to the late 1950's at MIT and is described in detail in Steven Levy's book **Hackers**. This book also includes a lot of information about Stallman himself and the development of Lisp, a programming language developed at MIT that underlies Emacs.)

## 1.1 Through Version 18

Although the history of the early versions of GNU Emacs is unclear, the history is well-known from the middle of 1985. A time line is:

- GNU Emacs version 15 (15.34) was released sometime in 1984 or 1985 and shared some code with a version of Emacs written by James Gosling (the same James Gosling who later created the Java language).
- GNU Emacs version 16 (first released version was 16.56) was released on July 15, 1985. All Gosling code was removed due to potential copyright problems with the code.
- version 16.57: released on September 16, 1985.
- versions 16.58, 16.59: released on September 17, 1985.
- version 16.60: released on September 19, 1985. These later version 16's incorporated patches from the net, esp. for getting Emacs to work under System V.
- version 17.36 (first official v17 release) released on December 20, 1985. Included a TeX-able user manual. First official unpatched version that worked on vanilla System V machines.
- version 17.43 (second official v17 release) released on January 25, 1986.
- version 17.45 released on January 30, 1986.
- version 17.46 released on February 4, 1986.
- version 17.48 released on February 10, 1986.
- version 17.49 released on February 12, 1986.
- version 17.55 released on March 18, 1986.
- version 17.57 released on March 27, 1986.
- version 17.58 released on April 4, 1986.
- version 17.61 released on April 12, 1986.
- version 17.63 released on May 7, 1986.
- version 17.64 released on May 12, 1986.

- version 18.24 (a beta version) released on October 2, 1986.
- version 18.30 (a beta version) released on November 15, 1986.
- version 18.31 (a beta version) released on November 23, 1986.
- version 18.32 (a beta version) released on December 7, 1986.
- version 18.33 (a beta version) released on December 12, 1986.
- version 18.35 (a beta version) released on January 5, 1987.
- version 18.36 (a beta version) released on January 21, 1987.
- January 27, 1987: The Great Usenet Renaming. net.emacs is now comp.emacs.
- version 18.37 (a beta version) released on February 12, 1987.
- version 18.38 (a beta version) released on March 3, 1987.
- version 18.39 (a beta version) released on March 14, 1987.
- version 18.40 (a beta version) released on March 18, 1987.
- version 18.41 (the first “official” release) released on March 22, 1987.
- version 18.45 released on June 2, 1987.
- version 18.46 released on June 9, 1987.
- version 18.47 released on June 18, 1987.
- version 18.48 released on September 3, 1987.
- version 18.49 released on September 18, 1987.
- version 18.50 released on February 13, 1988.
- version 18.51 released on May 7, 1988.
- version 18.52 released on September 1, 1988.
- version 18.53 released on February 24, 1989.
- version 18.54 released on April 26, 1989.
- version 18.55 released on August 23, 1989. This is the earliest version that is still available by FTP.
- version 18.56 released on January 17, 1991.
- version 18.57 released late January, 1991.
- version 18.58 released ?????.
- version 18.59 released October 31, 1992.

## 1.2 Lucid Emacs

Lucid Emacs was developed by the (now-defunct) Lucid Inc., a maker of C++ and Lisp development environments. It began when Lucid decided they wanted to use Emacs as the editor and cornerstone of their C++ development environment (called “Energize”). They needed many features that were not available in the existing version of GNU Emacs (version 18.5something), in particular good and integrated support for GUI elements such as mouse support, multiple fonts, multiple window-system windows, etc. A branch of GNU Emacs called Epoch, written at the University of Illinois, existed that supplied many of these features; however, Lucid needed more than what existed in Epoch. At the time, the Free Software Foundation was working on version 19 of Emacs (this was sometime around 1991),

which was planned to have similar features, and so Lucid decided to work with the Free Software Foundation. Their plan was to add features that they needed, and coordinate with the FSF so that the features would get included back into Emacs version 19.

Delays in the release of version 19 occurred, however (resulting in it finally being released more than a year after what was initially planned), and Lucid encountered unexpected technical resistance in getting their changes merged back into version 19, so they decided to release their own version of Emacs, which became Lucid Emacs 19.0.

The initial authors of Lucid Emacs were Matthieu Devin, Harlan Sexton, and Eric Benson, and the work was later taken over by Jamie Zawinski, who became “Mr. Lucid Emacs” for many releases.

A time line for Lucid Emacs is

- version 19.0 shipped with Energize 1.0, April 1992.
- version 19.1 released June 4, 1992.
- version 19.2 released June 19, 1992.
- version 19.3 released September 9, 1992.
- version 19.4 released January 21, 1993.
- version 19.5 was a repackaging of 19.4 with a few bug fixes and shipped with Energize 2.0. Never released to the net.
- version 19.6 released April 9, 1993.
- version 19.7 was a repackaging of 19.6 with a few bug fixes and shipped with Energize 2.1. Never released to the net.
- version 19.8 released September 6, 1993.
- version 19.9 released January 12, 1994.
- version 19.10 released May 27, 1994.

### 1.3 GNU Emacs 19

About a year after the initial release of Lucid Emacs, the FSF released a beta of their version of Emacs 19 (referred to here as “GNU Emacs”). By this time, the current version of Lucid Emacs was 19.6. (Strangely, the first released beta from the FSF was GNU Emacs 19.7.) A time line for GNU Emacs version 19 is

- version 19.8 (beta) released May 27, 1993.
- version 19.9 (beta) released May 27, 1993.
- version 19.10 (beta) released May 30, 1993.
- version 19.11 (beta) released June 1, 1993.
- version 19.12 (beta) released June 2, 1993.
- version 19.13 (beta) released June 8, 1993.
- version 19.14 (beta) released June 17, 1993.
- version 19.15 (beta) released June 19, 1993.
- version 19.16 (beta) released July 6, 1993.
- version 19.17 (beta) released late July, 1993.
- version 19.18 (beta) released August 9, 1993.

- version 19.19 (beta) released August 15, 1993.
- version 19.20 (beta) released November 17, 1993.
- version 19.21 (beta) released November 17, 1993.
- version 19.22 (beta) released November 28, 1993.
- version 19.23 (beta) released May 17, 1994.
- version 19.24 (beta) released May 16, 1994.
- version 19.25 (beta) released June 3, 1994.
- version 19.26 (beta) released September 11, 1994.
- version 19.27 (beta) released September 14, 1994.
- version 19.28 (first “official” release) released November 1, 1994.
- version 19.29 released June 21, 1995.
- version 19.30 released November 24, 1995.
- version 19.31 released May 25, 1996.
- version 19.32 released July 31, 1996.
- version 19.33 released August 11, 1996.
- version 19.34 released August 21, 1996.
- version 19.34b released September 6, 1996.

In some ways, GNU Emacs 19 was better than Lucid Emacs; in some ways, worse. Lucid soon began incorporating features from GNU Emacs 19 into Lucid Emacs; the work was mostly done by Richard Mlynarik, who had been working on and using GNU Emacs for a long time (back as far as version 16 or 17).

## 1.4 GNU Emacs 20

On February 2, 1997 work began on GNU Emacs to integrate Mule. The first release was made in September of that year.

A timeline for Emacs 20 is

- version 20.1 released September 17, 1997.
- version 20.2 released September 20, 1997.
- version 20.3 released August 19, 1998.

## 1.5 XEmacs

Around the time that Lucid was developing Energize, Sun Microsystems was developing their own development environment (called “SPARCWorks”) and also decided to use Emacs. They joined forces with the Epoch team at the University of Illinois and later with Lucid. The maintainer of the last-released version of Epoch was Marc Andreessen, but he dropped out and the Epoch project, headed by Simon Kaplan, lured Chuck Thompson away from a system administration job to become the primary Lucid Emacs author for Epoch and Sun. Chuck’s area of specialty became the redisplay engine (he replaced the old Lucid Emacs redisplay engine with a ported version from Epoch and then later rewrote it from scratch). Sun also hired Ben Wing (the author of Win-Emacs, a port of Lucid Emacs to Microsoft Windows 3.1) in 1993, for what was initially a one-month contract to fix some

event problems but later became a many-year involvement, punctuated by a six-month contract with Amdahl Corporation.

In 1994, Sun and Lucid agreed to rename Lucid Emacs to XEmacs (a name not favorable to either company); the first release called XEmacs was version 19.11. In June 1994, Lucid folded and Jamie quit to work for the newly formed Mosaic Communications Corp., later Netscape Communications Corp. (co-founded by the same Marc Andreessen, who had quit his Epoch job to work on a graphical browser for the World Wide Web). Chuck then became the primary maintainer of XEmacs, and put out versions 19.11 through 19.14 in conjunction with Ben. For 19.12 and 19.13, Chuck added the new redisplay and many other display improvements and Ben added MULE support (support for Asian and other languages) and redesigned most of the internal Lisp subsystems to better support the MULE work and the various other features being added to XEmacs. After 19.14 Chuck retired as primary maintainer and Steve Baur stepped in.

Soon after 19.13 was released, work began in earnest on the MULE internationalization code and the source tree was divided into two development paths. The MULE version was initially called 19.20, but was soon renamed to 20.0. In 1996 Martin Buchholz of Sun Microsystems took over the care and feeding of it and worked on it in parallel with the 19.14 development that was occurring at the same time. After much work by Martin, it was decided to release 20.0 ahead of 19.15 in February 1997. The source tree remained divided until 20.2 when the version 19 source was finally retired at version 19.16.

In 1997, Sun finally dropped all pretense of support for XEmacs and Martin Buchholz left the company in November. Since then, and mostly for the previous year, because Steve Baur was never paid to work on XEmacs, XEmacs has existed solely on the contributions of volunteers from the Free Software Community. Starting from 1997, Hrvoje Niksic and Kyle Jones have figured prominently in XEmacs development.

Many attempts have been made to merge XEmacs and GNU Emacs, but they have consistently failed.

A more detailed history is contained in the XEmacs About page.

A time line for XEmacs is

- version 19.11 (first XEmacs) released September 13, 1994.
- version 19.12 released June 23, 1995.
- version 19.13 released September 1, 1995.
- version 19.14 released June 23, 1996.
- version 20.0 released February 9, 1997.
- version 19.15 released March 28, 1997.
- version 20.1 (not released to the net) April 15, 1997.
- version 20.2 released May 16, 1997.
- version 19.16 released October 31, 1997.
- version 20.3 (the first stable version of XEmacs 20.x) released November 30, 1997.
- version 20.4 released February 28, 1998.
- version 21.0.60 released December 10, 1998. (The version naming scheme was changed at this point: [a] the second version number is odd for stable versions, even for beta versions; [b] a third version number is added, replacing the "beta xxx" ending for beta

versions and allowing for periodic maintenance releases for stable versions. Therefore, 21.0 was never "officially" released; similarly for 21.2, etc.)

- version 21.0.61 released January 4, 1999.
- version 21.0.63 released February 3, 1999.
- version 21.0.64 released March 1, 1999.
- version 21.0.65 released March 5, 1999.
- version 21.0.66 released March 12, 1999.
- version 21.0.67 released March 25, 1999.
- version 21.1.2 released May 14, 1999. (This is the followup to 21.0.67. The second version number was bumped to indicate the beginning of the "stable" series.)
- version 21.1.3 released June 26, 1999.
- version 21.1.4 released July 8, 1999.
- version 21.1.6 released August 14, 1999. (There was no 21.1.5.)
- version 21.1.7 released September 26, 1999.
- version 21.1.8 released November 2, 1999.
- version 21.1.9 released February 13, 2000.
- version 21.1.10 released May 7, 2000.
- version 21.1.10a released June 24, 2000.
- version 21.1.11 released July 18, 2000.
- version 21.1.12 released August 5, 2000.
- version 21.1.13 released January 7, 2001.
- version 21.1.14 released January 27, 2001.
- version 21.2.9 released February 3, 1999.
- version 21.2.10 released February 5, 1999.
- version 21.2.11 released March 1, 1999.
- version 21.2.12 released March 5, 1999.
- version 21.2.13 released March 12, 1999.
- version 21.2.14 released May 14, 1999.
- version 21.2.15 released June 4, 1999.
- version 21.2.16 released June 11, 1999.
- version 21.2.17 released June 22, 1999.
- version 21.2.18 released July 14, 1999.
- version 21.2.19 released July 30, 1999.
- version 21.2.20 released November 10, 1999.
- version 21.2.21 released November 28, 1999.
- version 21.2.22 released November 29, 1999.
- version 21.2.23 released December 7, 1999.
- version 21.2.24 released December 14, 1999.
- version 21.2.25 released December 24, 1999.

- version 21.2.26 released December 31, 1999.
- version 21.2.27 released January 18, 2000.
- version 21.2.28 released February 7, 2000.
- version 21.2.29 released February 16, 2000.
- version 21.2.30 released February 21, 2000.
- version 21.2.31 released February 23, 2000.
- version 21.2.32 released March 20, 2000.
- version 21.2.33 released May 1, 2000.
- version 21.2.34 released May 28, 2000.
- version 21.2.35 released July 19, 2000.
- version 21.2.36 released October 4, 2000.
- version 21.2.37 released November 14, 2000.
- version 21.2.38 released December 5, 2000.
- version 21.2.39 released December 31, 2000.
- version 21.2.40 released January 8, 2001.
- version 21.2.41 released January 17, 2001.
- version 21.2.42 released January 20, 2001.
- version 21.2.43 released January 26, 2001.
- version 21.2.44 released February 8, 2001.
- version 21.2.45 released February 23, 2001.
- version 21.2.46 released March 21, 2001.



## 2 XEmacs From the Outside

XEmacs appears to the outside world as an editor, but it is really a Lisp environment. At its heart is a Lisp interpreter; it also “happens” to contain many specialized object types (e.g. buffers, windows, frames, events) that are useful for implementing an editor. Some of these objects (in particular windows and frames) have displayable representations, and XEmacs provides a function `redisplay()` that ensures that the display of all such objects matches their internal state. Most of the time, a standard Lisp environment is in a **read-eval-print** loop—i.e. “read some Lisp code, execute it, and print the results”. XEmacs has a similar loop:

- read an event
- dispatch the event (i.e. “do it”)
- redisplay

Reading an event is done using the Lisp function `next-event`, which waits for something to happen (typically, the user presses a key or moves the mouse) and returns an event object describing this. Dispatching an event is done using the Lisp function `dispatch-event`, which looks up the event in a keymap object (a particular kind of object that associates an event with a Lisp function) and calls that function. The function “does” what the user has requested by changing the state of particular frame objects, buffer objects, etc. Finally, `redisplay()` is called, which updates the display to reflect those changes just made. Thus is an “editor” born.

Note that you do not have to use XEmacs as an editor; you could just as well make it do your taxes, compute pi, play bridge, etc. You’d just have to write functions to do those operations in Lisp.



### 3 The Lisp Language

Lisp is a general-purpose language that is higher-level than C and in many ways more powerful than C. Powerful dialects of Lisp such as Common Lisp are probably much better languages for writing very large applications than is C. (Unfortunately, for many non-technical reasons C and its successor C++ have become the dominant languages for application development. These languages are both inadequate for extremely large applications, which is evidenced by the fact that newer, larger programs are becoming ever harder to write and are requiring ever more programmers despite great increases in C development environments; and by the fact that, although hardware speeds and reliability have been growing at an exponential rate, most software is still generally considered to be slow and buggy.)

The new Java language holds promise as a better general-purpose development language than C. Java has many features in common with Lisp that are not shared by C (this is not a coincidence, since Java was designed by James Gosling, a former Lisp hacker). This will be discussed more later.

For those used to C, here is a summary of the basic differences between C and Lisp:

1. Lisp has an extremely regular syntax. Every function, expression, and control statement is written in the form

```
(func arg1 arg2 ...)
```

This is as opposed to C, which writes functions as

```
func(arg1, arg2, ...)
```

but writes expressions involving operators as (e.g.)

```
arg1 + arg2
```

and writes control statements as (e.g.)

```
while (expr) { statement1 ; statement2 ; ... }
```

Lisp equivalents of the latter two would be

```
(+ arg1 arg2 ...)
```

and

```
(while expr statement1 statement2 ...)
```

2. Lisp is a safe language. Assuming there are no bugs in the Lisp interpreter/compiler, it is impossible to write a program that “core dumps” or otherwise causes the machine to execute an illegal instruction. This is very different from C, where perhaps the most common outcome of a bug is exactly such a crash. A corollary of this is that the C operation of casting a pointer is impossible (and unnecessary) in Lisp, and that it is impossible to access memory outside the bounds of an array.
3. Programs and data are written in the same form. The parenthesis-enclosing form described above for statements is the same form used for the most common data type in Lisp, the list. Thus, it is possible to represent any Lisp program using Lisp data types, and for one program to construct Lisp statements and then dynamically **evaluate** them, or cause them to execute.
4. All objects are **dynamically typed**. This means that part of every object is an indication of what type it is. A Lisp program can manipulate an object without knowing what type it is, and can query an object to determine its type. This means that, correspondingly,

variables and function parameters can hold objects of any type and are not normally declared as being of any particular type. This is opposed to the **static typing** of C, where variables can hold exactly one type of object and must be declared as such, and objects do not contain an indication of their type because it's implicit in the variables they are stored in. It is possible in C to have a variable hold different types of objects (e.g. through the use of `void *` pointers or variable-argument functions), but the type information must then be passed explicitly in some other fashion, leading to additional program complexity.

5. Allocated memory is automatically reclaimed when it is no longer in use. This operation is called **garbage collection** and involves looking through all variables to see what memory is being pointed to, and reclaiming any memory that is not pointed to and is thus “inaccessible” and out of use. This is as opposed to C, in which allocated memory must be explicitly reclaimed using `free()`. If you simply drop all pointers to memory without freeing it, it becomes “leaked” memory that still takes up space. Over a long period of time, this can cause your program to grow and grow until it runs out of memory.
6. Lisp has built-in facilities for handling errors and exceptions. In C, when an error occurs, usually either the program exits entirely or the routine in which the error occurs returns a value indicating this. If an error occurs in a deeply-nested routine, then every routine currently called must unwind itself normally and return an error value back up to the next routine. This means that every routine must explicitly check for an error in all the routines it calls; if it does not do so, unexpected and often random behavior results. This is an extremely common source of bugs in C programs. An alternative would be to do a non-local exit using `longjmp()`, but that is often very dangerous because the routines that were exited past had no opportunity to clean up after themselves and may leave things in an inconsistent state, causing a crash shortly afterwards.

Lisp provides mechanisms to make such non-local exits safe. When an error occurs, a routine simply signals that an error of a particular class has occurred, and a non-local exit takes place. Any routine can trap errors occurring in routines it calls by registering an error handler for some or all classes of errors. (If no handler is registered, a default handler, generally installed by the top-level event loop, is executed; this prints out the error and continues.) Routines can also specify cleanup code (called an **unwind-protect**) that will be called when control exits from a block of code, no matter how that exit occurs—i.e. even if a function deeply nested below it causes a non-local exit back to the top level.

Note that this facility has appeared in some recent vintages of C, in particular Visual C++ and other PC compilers written for the Microsoft Win32 API.

7. In Emacs Lisp, local variables are **dynamically scoped**. This means that if you declare a local variable in a particular function, and then call another function, that subfunction can “see” the local variable you declared. This is actually considered a bug in Emacs Lisp and in all other early dialects of Lisp, and was corrected in Common Lisp. (In Common Lisp, you can still declare dynamically scoped variables if you want to—they are sometimes useful—but variables by default are **lexically scoped** as in C.)

For those familiar with Lisp, Emacs Lisp is modelled after MacLisp, an early dialect of Lisp developed at MIT (no relation to the Macintosh computer). There is a Common Lisp compatibility package available for Emacs that provides many of the features of Common Lisp.

The Java language is derived in many ways from C, and shares a similar syntax, but has the following features in common with Lisp (and different from C):

1. Java is a safe language, like Lisp.
2. Java provides garbage collection, like Lisp.
3. Java has built-in facilities for handling errors and exceptions, like Lisp.
4. Java has a type system that combines the best advantages of both static and dynamic typing. Objects (except very simple types) are explicitly marked with their type, as in dynamic typing; but there is a hierarchy of types and functions are declared to accept only certain types, thus providing the increased compile-time error-checking of static typing.

The Java language also has some negative attributes:

1. Java uses the edit/compile/run model of software development. This makes it hard to use interactively. For example, to use Java like `bc` it is necessary to write a special purpose, albeit tiny, application. In Emacs Lisp, a calculator comes built-in without any effort - one can always just type an expression in the `*scratch*` buffer.
2. Java tries too hard to enforce, not merely enable, portability, making ordinary access to standard OS facilities painful. Java has an **agenda** I think this is why `chdir` is not part of standard Java, which is inexcusable.

Unfortunately, there is no perfect language. Static typing allows a compiler to catch programmer errors and produce more efficient code, but makes programming more tedious and less fun. For the foreseeable future, an Ideal Editing and Programming Environment (and that is what XEmacs aspires to) will be programmable in multiple languages: high level ones like Lisp for user customization and prototyping, and lower level ones for infrastructure and industrial strength applications. If I had my way, XEmacs would be friendly towards the Python, Scheme, C++, ML, etc... communities. But there are serious technical difficulties to achieving that goal.

The word **application** in the previous paragraph was used intentionally. XEmacs implements an API for programs written in Lisp that makes it a full-fledged application platform, very much like an OS inside the real OS.



## 4 XEmacs From the Perspective of Building

The heart of XEmacs is the Lisp environment, which is written in C. This is contained in the `'src/'` subdirectory. Underneath `'src/'` are two subdirectories of header files: `'s/'` (header files for particular operating systems) and `'m/'` (header files for particular machine types). In practice the distinction between the two types of header files is blurred. These header files define or undefine certain preprocessor constants and macros to indicate particular characteristics of the associated machine or operating system. As part of the configure process, one `'s/'` file and one `'m/'` file is identified for the particular environment in which XEmacs is being built.

XEmacs also contains a great deal of Lisp code. This implements the operations that make XEmacs useful as an editor as well as just a Lisp environment, and also contains many add-on packages that allow XEmacs to browse directories, act as a mail and Usenet news reader, compile Lisp code, etc. There is actually more Lisp code than C code associated with XEmacs, but much of the Lisp code is peripheral to the actual operation of the editor. The Lisp code all lies in subdirectories underneath the `'lisp/'` directory.

The `'lwlib/'` directory contains C code that implements a generalized interface onto different X widget toolkits and also implements some widgets of its own that behave like Motif widgets but are faster, free, and in some cases more powerful. The code in this directory compiles into a library and is mostly independent from XEmacs.

The `'etc/'` directory contains various data files associated with XEmacs. Some of them are actually read by XEmacs at startup; others merely contain useful information of various sorts.

The `'lib-src/'` directory contains C code for various auxiliary programs that are used in connection with XEmacs. Some of them are used during the build process; others are used to perform certain functions that cannot conveniently be placed in the XEmacs executable (e.g. the `'movemail'` program for fetching mail out of `'/var/spool/mail'`, which must be setgid to `'mail'` on many systems; and the `'gnuclient'` program, which allows an external script to communicate with a running XEmacs process).

The `'man/'` directory contains the sources for the XEmacs documentation. It is mostly in a form called Texinfo, which can be converted into either a printed document (by passing it through `TEX`) or into on-line documentation called `info les`.

The `'info/'` directory contains the results of formatting the XEmacs documentation as `info les`, for on-line use. These files are used when you enter the Info system using `C-h i` or through the Help menu.

The `'dynodump/'` directory contains auxiliary code used to build XEmacs on Solaris platforms.

The other directories contain various miscellaneous code and information that is not normally used or needed.

The first step of building involves running the `'configure'` program and passing it various parameters to specify any optional features you want and compiler arguments and such, as described in the `'INSTALL'` file. This determines what the build environment is, chooses the appropriate `'s/'` and `'m/'` file, and runs a series of tests to determine many details about your environment, such as which library functions are available and exactly how they work. The reason for running these tests is that it allows XEmacs to be compiled on a much wider

variety of platforms than those that the XEmacs developers happen to be familiar with, including various sorts of hybrid platforms. This is especially important now that many operating systems give you a great deal of control over exactly what features you want installed, and allow for easy upgrading of parts of a system without upgrading the rest. It would be impossible to pre-determine and pre-specify the information for all possible configurations.

In fact, the `'s/` and `'m/` files are basically *evil*, since they contain unmaintainable platform-specific hard-coded information. XEmacs has been moving in the direction of having all system-specific information be determined dynamically by `'configure`'. Perhaps someday we can `rm -rf src/s src/m`.

When `configure` is done running, it generates `'Makefile`'s and `'GNUmakefile`'s and the file `'src/config.h`' (which describes the features of your system) from template files. You then run `'make`', which compiles the auxiliary code and programs in `'lib-src/` and `'lwlib/` and the main XEmacs executable in `'src/`'. The result of compiling and linking is an executable called `'temacs`', which is *not* the final XEmacs executable. `'temacs`' by itself is not intended to function as an editor or even display any windows on the screen, and if you simply run it, it will exit immediately. The `'Makefile`' runs `'temacs`' with certain options that cause it to initialize itself, read in a number of basic Lisp files, and then dump itself out into a new executable called `'xemacs`'. This new executable has been pre-initialized and contains pre-digested Lisp code that is necessary for the editor to function (this includes most basic editing functions, e.g. `kill-line`, that can be defined in terms of other Lisp primitives; some initialization code that is called when certain objects, such as frames, are created; and all of the standard keybindings and code for the actions they result in). This executable, `'xemacs`', is the executable that you run to use the XEmacs editor.

Although `'temacs`' is not intended to be run as an editor, it can, by using the incantation `temacs -batch -l loadup.el run-temacs`. This is useful when the dumping procedure described above is broken, or when using certain program debugging tools such as Purify. These tools get mighty confused by the tricks played by the XEmacs build process, such as allocating memory in one process, and freeing it in the next.

## 5 Build-Time Dependencies

This is a collection of random notes on build-time dependencies as of about XEmacs 21.5.11. Of course we use ‘make’ to manage most dependencies, especially for the C code. The main thing here is for the Release Engineer to run the ‘src/make-src-depend’ script every so often, at least at every release.

However, since most of XEmacs is written in Lisp, and we compile and preload the Lisp for efficiency, managing Lisp compilation using ‘make’ would imply running XEmacs hundreds of times. This would make the build process unbearably long. Thus those processes that require running the same Lisp programs on many files are managed using Lisp driver functions rather than ‘make’. The situation is further complicated by the fact that documentation strings are kept in an external database, and referenced in the dumped XEmacs by file offset. Finally, the Lisp files are processed to collect autoloaded function information and customize dependencies, which are then written into generated Lisp files.

About this, Ben sez:

1. Redumping depends on up-to-date dumped ‘.elc’ files and ‘DOC’ but not directly on auto-autoloads.
2. Rebuilding dumped ‘.elc’ files depends on auto-autoloads being up-to-date.
3. Building the ‘DOC’ file depends on up-to-date dumped ‘.elc’ files but not directly on auto-autoloads.
4. Recompiling anything depends on ‘bytecomp.elc’ and ‘byte-optimize.elc’ being up-to-date.

Put these together and you’ll see it’s perfectly acceptable to build auto-autoloads *after* dumping if no ‘.elc’ files are out-of-date.

These Lisp driver programs typically run from `temacs`, not a dumped XEmacs. The simplest (but time-consuming) way to achieve a sane environment for running Lisp is to load ‘loadup.el’ or ‘loadup-el.el’. (The latter is used to avoid loading possibly out-of-date compiled Lisp files.) If this is not done, you have to construct the environment yourself. See ‘dumped-lisp.el’ to see how it is done in the dumped XEmacs.

One potential gotcha is that very early customizations are now handled by adding the definitions to the special variable `custom-declare-variable-list`, defined in ‘subr.el’. If you use any higher-level functionality that might load ‘custom.el’, but you do not need ‘subr.el’, you should ‘defvar’ `custom-declare-variable-list` to prevent the ‘void-variable’ error. (Currently this is only needed for ‘make-docfile.el’.)



## 6 XEmacs From the Inside

Internally, XEmacs is quite complex, and can be very confusing. To simplify things, it can be useful to think of XEmacs as containing an event loop that “drives” everything, and a number of other subsystems, such as a Lisp engine and a redisplay mechanism. Each of these other subsystems exists simultaneously in XEmacs, and each has a certain state. The flow of control continually passes in and out of these different subsystems in the course of normal operation of the editor.

It is important to keep in mind that, most of the time, the editor is “driven” by the event loop. Except during initialization and batch mode, all subsystems are entered directly or indirectly through the event loop, and ultimately, control exits out of all subsystems back up to the event loop. This cycle of entering a subsystem, exiting back out to the event loop, and starting another iteration of the event loop occurs once each keystroke, mouse motion, etc.

If you’re trying to understand a particular subsystem (other than the event loop), think of it as a “daemon” process or “servant” that is responsible for one particular aspect of a larger system, and periodically receives commands or environment changes that cause it to do something. Ultimately, these commands and environment changes are always triggered by the event loop. For example:

- The window and frame mechanism is responsible for keeping track of what windows and frames exist, what buffers are in them, etc. It is periodically given commands (usually from the user) to make a change to the current window/frame state: i.e. create a new frame, delete a window, etc.
- The buffer mechanism is responsible for keeping track of what buffers exist and what text is in them. It is periodically given commands (usually from the user) to insert or delete text, create a buffer, etc. When it receives a text-change command, it notifies the redisplay mechanism.
- The redisplay mechanism is responsible for making sure that windows and frames are displayed correctly. It is periodically told (by the event loop) to actually “do its job”, i.e. snoop around and see what the current state of the environment (mostly of the currently-existing windows, frames, and buffers) is, and make sure that state matches what’s actually displayed. It keeps lots and lots of information around (such as what is actually being displayed currently, and what the environment was last time it checked) so that it can minimize the work it has to do. It is also helped along in that whenever a relevant change to the environment occurs, the redisplay mechanism is told about this, so it has a pretty good idea of where it has to look to find possible changes and doesn’t have to look everywhere.
- The Lisp engine is responsible for executing the Lisp code in which most user commands are written. It is entered through a call to `eval` or `funcall`, which occurs as a result of dispatching an event from the event loop. The functions it calls issue commands to the buffer mechanism, the window/frame subsystem, etc.
- The Lisp allocation subsystem is responsible for keeping track of Lisp objects. It is given commands from the Lisp engine to allocate objects, garbage collect, etc.

etc.

The important idea here is that there are a number of independent subsystems each with its own responsibility and persistent state, just like different employees in a company, and each subsystem is periodically given commands from other subsystems. Commands can flow from any one subsystem to any other, but there is usually some sort of hierarchy, with all commands originating from the event subsystem.

XEmacs is entered in `main()`, which is in `'emacs.c'`. When this is called the first time (in a properly-invoked `'temacs'`), it does the following:

1. It does some very basic environment initializations, such as determining where it and its directories (e.g. `'lisp/'` and `'etc/'`) reside and setting up signal handlers.
2. It initializes the entire Lisp interpreter.
3. It sets the initial values of many built-in variables (including many variables that are visible to Lisp programs), such as the global keymap object and the built-in faces (a face is an object that describes the display characteristics of text). This involves creating Lisp objects and thus is dependent on step (2).
4. It performs various other initializations that are relevant to the particular environment it is running in, such as retrieving environment variables, determining the current date and the user who is running the program, examining its standard input, creating any necessary file descriptors, etc.
5. At this point, the C initialization is complete. A Lisp program that was specified on the command line (usually `'loadup.el'`) is called (temacs is normally invoked as `temacs -batch -l loadup.el dump`). `'loadup.el'` loads all of the other Lisp files that are needed for the operation of the editor, calls the `dump-emacs` function to write out `'xemacs'`, and then kills the temacs process.

When `'xemacs'` is then run, it only redoes steps (1) and (4) above; all variables already contain the values they were set to when the executable was dumped, and all memory that was allocated with `malloc()` is still around. (XEmacs knows whether it is being run as `'xemacs'` or `'temacs'` because it sets the global variable `initialized` to 1 after step (4) above.) At this point, `'xemacs'` calls a Lisp function to do any further initialization, which includes parsing the command-line (the C code can only do limited command-line parsing, which includes looking for the `'-batch'` and `'-l'` flags and a few other flags that it needs to know about before initialization is complete), creating the first frame (or `window` in standard window-system parlance), running the user's init file (usually the file `'emacs'` in the user's home directory), etc. The function to do this is usually called `normal-top-level`; `'loadup.el'` tells the C code about this function by setting its name as the value of the Lisp variable `top-level`.

When the Lisp initialization code is done, the C code enters the event loop, and stays there for the duration of the XEmacs process. The code for the event loop is contained in `'cmdloop.c'`, and is called `Fcommand_loop_1()`. Note that this event loop could very well be written in Lisp, and in fact a Lisp version exists; but apparently, doing this makes XEmacs run noticeably slower.

Notice how much of the initialization is done in Lisp, not in C. In general, XEmacs tries to move as much code as is possible into Lisp. Code that remains in C is code that implements the Lisp interpreter itself, or code that needs to be very fast, or code that needs to do system calls or other such stuff that needs to be done in C, or code that needs to have access to “forbidden” structures. (One conscious aspect of the design of Lisp under

XEmacs is a clean separation between the external interface to a Lisp object's functionality and its internal implementation. Part of this design is that Lisp programs are forbidden from accessing the contents of the object other than through using a standard API. In this respect, XEmacs Lisp is similar to modern Lisp dialects but differs from GNU Emacs, which tends to expose the implementation and allow Lisp programs to look at it directly. The major advantage of hiding the implementation is that it allows the implementation to be redesigned without affecting any Lisp programs, including those that might want to be "clever" by looking directly at the object's contents and possibly manipulating them.)

Moving code into Lisp makes the code easier to debug and maintain and makes it much easier for people who are not XEmacs developers to customize XEmacs, because they can make a change with much less chance of obscure and unwanted interactions occurring than if they were to change the C code.



## 7 The XEmacs Object System (Abstractly Speaking)

At the heart of the Lisp interpreter is its management of objects. XEmacs Lisp contains many built-in objects, some of which are simple and others of which can be very complex; and some of which are very common, and others of which are rarely used or are only used internally. (Since the Lisp allocation system, with its automatic reclamation of unused storage, is so much more convenient than `malloc()` and `free()`, the C code makes extensive use of it in its internal operations.)

The basic Lisp objects are

<code>integer</code>	28 or 31 bits of precision, or 60 or 63 bits on 64-bit machines; the reason for this is described below when the internal Lisp object representation is described.
<code>float</code>	Same precision as a double in C.
<code>cons</code>	A simple container for two Lisp objects, used to implement lists and most other data structures in Lisp.
<code>char</code>	An object representing a single character of text; chars behave like integers in many ways but are logically considered text rather than numbers and have a different read syntax. (the read syntax for a char contains the char itself or some textual encoding of it—for example, a Japanese Kanji character might be encoded as ‘ <code>^[\$(B#&amp;^[ (B</code> ’ using the ISO-2022 encoding standard—rather than the numerical representation of the char; this way, if the mapping between chars and integers changes, which is quite possible for Kanji characters and other extended characters, the same character will still be created. Note that some primitives confuse chars and integers. The worst culprit is <code>eq</code> , which makes a special exception and considers a char to be <code>eq</code> to its integer equivalent, even though in no other case are objects of two different types <code>eq</code> . The reason for this monstrosity is compatibility with existing code; the separation of char from integer came fairly recently.)
<code>symbol</code>	An object that contains Lisp objects and is referred to by name; symbols are used to implement variables and named functions and to provide the equivalent of preprocessor constants in C.
<code>vector</code>	A one-dimensional array of Lisp objects providing constant-time access to any of the objects; access to an arbitrary object in a vector is faster than for lists, but the operations that can be done on a vector are more limited.
<code>string</code>	Self-explanatory; behaves much like a vector of chars but has a different read syntax and is stored and manipulated more compactly.
<code>bit-vector</code>	A vector of bits; similar to a string in spirit.
<code>compiled-function</code>	An object containing compiled Lisp code, known as <b>byte code</b>
<code>subr</code>	A Lisp primitive, i.e. a Lisp-callable function implemented in C.

Note that there is no basic “function” type, as in more powerful versions of Lisp (where it’s called a **closure**). XEmacs Lisp does not provide the closure semantics implemented

by Common Lisp and Scheme. The guts of a function in XEmacs Lisp are represented in one of four ways: a symbol specifying another function (when one function is an alias for another), a list (whose first element must be the symbol `lambda`) containing the function's source code, a compiled-function object, or a subr object. (In other words, given a symbol specifying the name of a function, calling `symbol-function` to retrieve the contents of the symbol's function cell will return one of these types of objects.)

XEmacs Lisp also contains numerous specialized objects used to implement the editor:

<b>buffer</b>	Stores text like a string, but is optimized for insertion and deletion and has certain other properties that can be set.
<b>frame</b>	An object with various properties whose displayable representation is a <b>window</b> in window-system parlance.
<b>window</b>	A section of a frame that displays the contents of a buffer; often called a <b>pane</b> in window-system parlance.
<b>window-configuration</b>	An object that represents a saved configuration of windows in a frame.
<b>device</b>	An object representing a screen on which frames can be displayed; equivalent to a <b>display</b> in the X Window System and a <b>TTY</b> in character mode.
<b>face</b>	An object specifying the appearance of text or graphics; it has properties such as font, foreground color, and background color.
<b>marker</b>	An object that refers to a particular position in a buffer and moves around as text is inserted and deleted to stay in the same relative position to the text around it.
<b>extent</b>	Similar to a marker but covers a range of text in a buffer; can also specify properties of the text, such as a face in which the text is to be displayed, whether the text is invisible or unmodifiable, etc.
<b>event</b>	Generated by calling <code>next-event</code> and contains information describing a particular event happening in the system, such as the user pressing a key or a process terminating.
<b>keymap</b>	An object that maps from events (described using lists, vectors, and symbols rather than with an event object because the mapping is for classes of events, rather than individual events) to functions to execute or other events to recursively look up; the functions are described by name, using a symbol, or using lists to specify the function's code.
<b>glyph</b>	An object that describes the appearance of an image (e.g. <code>pixmap</code> ) on the screen; glyphs can be attached to the beginning or end of extents and in some future version of XEmacs will be able to be inserted directly into a buffer.
<b>process</b>	An object that describes a connection to an externally-running process.

There are some other, less-commonly-encountered general objects:

<b>hash-table</b>	An object that maps from an arbitrary Lisp object to another arbitrary Lisp object, using hashing for fast lookup.
-------------------	--

**obarray** A limited form of hash-table that maps from strings to symbols; obarrays are used to look up a symbol given its name and are not actually their own object type but are kludgily represented using vectors with hidden fields (this representation derives from GNU Emacs).

**specifier** A complex object used to specify the value of a display property; a default value is given and different values can be specified for particular frames, buffers, windows, devices, or classes of device.

**char-table** An object that maps from chars or classes of chars to arbitrary Lisp objects; internally char tables use a complex nested-vector representation that is optimized to the way characters are represented as integers.

**range-table** An object that maps from ranges of integers to arbitrary Lisp objects.

And some strange special-purpose objects:

**charset**

**coding-system**

Objects used when MULE, or multi-lingual/Asian-language, support is enabled.

**color-instance**

**font-instance**

**image-instance**

An object that encapsulates a window-system resource; instances are mostly used internally but are exposed on the Lisp level for cleanness of the specifier model and because it's occasionally useful for Lisp program to create or query the properties of instances.

**subwindow**

An object that encapsulate a **subwindow** resource, i.e. a window-system child window that is drawn into by an external process; this object should be integrated into the glyph system but isn't yet, and may change form when this is done.

**tooltalk-message**

**tooltalk-pattern**

Objects that represent resources used in the ToolTalk interprocess communication protocol.

**toolbar-button**

An object used in conjunction with the toolbar.

And objects that are only used internally:

**opaque** A generic object for encapsulating arbitrary memory; this allows you the generality of `malloc()` and the convenience of the Lisp object system.

**lstream** A buffering I/O stream, used to provide a unified interface to anything that can accept output or provide input, such as a file descriptor, a stdio stream, a chunk of memory, a Lisp buffer, a Lisp string, etc.; it's a Lisp object to make its memory management more convenient.

**char-table-entry**

Subsidiary objects in the internal char-table representation.

**extent-auxiliary****menubar-data****toolbar-data**

Various special-purpose objects that are basically just used to encapsulate memory for particular subsystems, similar to the more general “opaque” object.

**symbol-value-forward****symbol-value-buffer-local****symbol-value-varalias****symbol-value-lisp-magic**

Special internal-only objects that are placed in the value cell of a symbol to indicate that there is something special with this variable – e.g. it has no value, it mirrors another variable, or it mirrors some C variable; there is really only one kind of object, called a **symbol-value-magic** but it is sort-of halfway kludged into semi-different object types.

Some types of objects are **permanent**, meaning that once created, they do not disappear until explicitly destroyed, using a function such as **delete-buffer**, **delete-window**, **delete-frame**, etc. Others will disappear once they are not longer used, through the garbage collection mechanism. Buffers, frames, windows, devices, and processes are among the objects that are permanent. Note that some objects can go both ways: Faces can be created either way; extents are normally permanent, but detached extents (extents not referring to any text, as happens to some extents when the text they are referring to is deleted) are temporary. Note that some permanent objects, such as faces and coding systems, cannot be deleted. Note also that windows are unique in that they can be *undeleted* after having previously been deleted. (This happens as a result of restoring a window configuration.)

Note that many types of objects have a **read syntax**, i.e. a way of specifying an object of that type in Lisp code. When you load a Lisp file, or type in code to be evaluated, what really happens is that the function **read** is called, which reads some text and creates an object based on the syntax of that text; then **eval** is called, which possibly does something special; then this loop repeats until there’s no more text to read. (**eval** only actually does something special with symbols, which causes the symbol’s value to be returned, similar to referencing a variable; and with conses [i.e. lists], which cause a function invocation. All other values are returned unchanged.)

The read syntax

17297

converts to an integer whose value is 17297.

1.983e-4

converts to a float whose value is 1.983e-4, or .0001983.

?b

converts to a char that represents the lowercase letter b.

?^[\$(B#&^[ (B

(where ‘^[’ actually is an ‘ESC’ character) converts to a particular Kanji character when using an ISO2022-based coding system for input. (To decode this goo: ‘ESC’ begins an

escape sequence; ‘ESC \$ (’ is a class of escape sequences meaning “switch to a 94x94 character set”; ‘ESC \$ ( B’ means “switch to Japanese Kanji”; ‘#’ and ‘&’ collectively index into a 94-by-94 array of characters [subtract 33 from the ASCII value of each character to get the corresponding index]; ‘ESC (’ is a class of escape sequences meaning “switch to a 94 character set”; ‘ESC (B’ means “switch to US ASCII”. It is a coincidence that the letter ‘B’ is used to denote both Japanese Kanji and US ASCII. If the first ‘B’ were replaced with an ‘A’, you’d be requesting a Chinese Hanzi character from the GB2312 character set.)

```
"foobar"
```

converts to a string.

```
foobar
```

converts to a symbol whose name is "foobar". This is done by looking up the string equivalent in the global variable `obarray`, whose contents should be an obarray. If no symbol is found, a new symbol with the name "foobar" is automatically created and added to `obarray`; this process is called **interning** the symbol.

```
(foo . bar)
```

converts to a cons cell containing the symbols `foo` and `bar`.

```
(1 a 2.5)
```

converts to a three-element list containing the specified objects (note that a list is actually a set of nested conses; see the XEmacs Lisp Reference).

```
[1 a 2.5]
```

converts to a three-element vector containing the specified objects.

```
#[... ..]
```

converts to a compiled-function object (the actual contents are not shown since they are not relevant here; look at a file that ends with ‘.elc’ for examples).

```
##01110110
```

converts to a bit-vector.

```
#s(hash-table ... ..)
```

converts to a hash table (the actual contents are not shown).

```
#s(range-table ... ..)
```

converts to a range table (the actual contents are not shown).

```
#s(char-table ... ..)
```

converts to a char table (the actual contents are not shown).

Note that the `#s()` syntax is the general syntax for structures, which are not really implemented in XEmacs Lisp but should be.

When an object is printed out (using `print` or a related function), the read syntax is used, so that the same object can be read in again.

The other objects do not have read syntaxes, usually because it does not really make sense to create them in this fashion (i.e. processes, where it doesn’t make sense to have a subprocess created as a side effect of reading some Lisp code), or because they can’t be created at all (e.g. subrs). Permanent objects, as a rule, do not have a read syntax; nor do most complex objects, which contain too much state to be easily initialized through a read syntax.



## 8 How Lisp Objects Are Represented in C

Lisp objects are represented in C using a 32-bit or 64-bit machine word (depending on the processor; i.e. DEC Alphas use 64-bit Lisp objects and most other processors use 32-bit Lisp objects). The representation stuffs a pointer together with a tag, as follows:

```
[ 3 3 2 2 2 2 2 2 2 2 2 1 1 1 1 1 1 1 1 1 0 0 0 0 0 0 0 0 0 0 ]
[ 1 0 9 8 7 6 5 4 3 2 1 0 9 8 7 6 5 4 3 2 1 0 9 8 7 6 5 4 3 2 1 0 ]
```

```
<-----> <->
      a pointer to a structure, or an integer      tag
```

A tag of 00 is used for all pointer object types, a tag of 10 is used for characters, and the other two tags 01 and 11 are joined together to form the integer object type. This representation gives us 31 bit integers and 30 bit characters, while pointers are represented directly without any bit masking or shifting. This representation, though, assumes that pointers to structs are always aligned to multiples of 4, so the lower 2 bits are always zero.

Lisp objects use the typedef `Lisp_Object`, but the actual C type used for the Lisp object can vary. It can be either a simple type (`long` on the DEC Alpha, `int` on other machines) or a structure whose fields are bit fields that line up properly (actually, a union of structures is used). Generally the simple integral type is preferable because it ensures that the compiler will actually use a machine word to represent the object (some compilers will use more general and less efficient code for unions and structs even if they can fit in a machine word). The union type, however, has the advantage of stricter type checking. If you accidentally pass an integer where a Lisp object is desired, you get a compile error. The choice of which type to use is determined by the preprocessor constant `USE_UNION_TYPE` which is defined via the `--use-union-type` option to `configure`.

Various macros are used to convert between Lisp\_Objects and the corresponding C type. Macros of the form `XINT()`, `XCHAR()`, `XSTRING()`, `XSYMBOL()`, do any required bit shifting and/or masking and cast it to the appropriate type. `XINT()` needs to be a bit tricky so that negative numbers are properly sign-extended. Since integers are stored left-shifted, if the right-shift operator does an arithmetic shift (i.e. it leaves the most-significant bit as-is rather than shifting in a zero, so that it mimics a divide-by-two even for negative numbers) the shift to remove the tag bit is enough. This is the case on all the systems we support.

Note that when `ERROR_CHECK_TYPECHECK` is defined, the converter macros become more complicated—they check the tag bits and/or the type field in the first four bytes of a record type to ensure that the object is really of the correct type. This is great for catching places where an incorrect type is being dereferenced—this typically results in a pointer being dereferenced as the wrong type of structure, with unpredictable (and sometimes not easily traceable) results.

There are similar `XSETTYPE()` macros that construct a Lisp object. These macros are of the form `XSETTYPE(lvalue, result)`, i.e. they have to be a statement rather than just used in an expression. The reason for this is that standard C doesn't let you "construct" a structure (but GCC does). Granted, this sometimes isn't too convenient; for the case of integers, at least, you can use the function `make_int()`, which constructs and *returns* an integer Lisp object. Note that the `XSETTYPE()` macros are also affected by `ERROR_CHECK_TYPECHECK` and make sure that the structure is of the right type in the case of record types, where the type is contained in the structure.

The C programmer is responsible for **guaranteeing** that a Lisp\_Object is the correct type before using the **XTYPE** macros. This is especially important in the case of lists. Use **XCAR** and **XCDR** if a Lisp\_Object is certainly a cons cell, else use **Fcar()** and **Fcdr()**. Trust other C code, but not Lisp code. On the other hand, if XEmacs has an internal logic error, it's better to crash immediately, so sprinkle **assert()**s and “unreachable” **abort()**s liberally about the source code. Where performance is an issue, use **type\_checking\_assert**, **bufpos\_checking\_assert**, and **gc\_checking\_assert**, which do nothing unless the corresponding configure error checking flag was specified.

## 9 Major Textual Changes

Sometimes major textual changes are made to the source. This means that a search-and-replace is done to change type names and such. Some people disagree with such changes, and certainly if done without good reason will just lead to headaches. But it's important to keep the code clean and understable, and consistent naming goes a long way towards this.

An example of the right way to do this was the so-called "great integral type renaming".

### 9.1 Great Integral Type Renaming

The purpose of this is to rationalize the names used for various integral types, so that they match their intended uses and follow consist conventions, and eliminate types that were not semantically different from each other.

The conventions are:

- All integral types that measure quantities of anything are signed. Some people disagree vociferously with this, but their arguments are mostly theoretical, and are vastly outweighed by the practical headaches of mixing signed and unsigned values, and more importantly by the far increased likelihood of inadvertent bugs: Because of the broken "viral" nature of unsigned quantities in C (operations involving mixed signed/unsigned are done unsigned, when exactly the opposite is nearly always wanted), even a single error in declaring a quantity unsigned that should be signed, or even the even more subtle error of comparing signed and unsigned values and forgetting the necessary cast, can be catastrophic, as comparisons will yield wrong results. `-Wsign-compare` is turned on specifically to catch this, but this tends to result in a great number of warnings when mixing signed and unsigned, and the casts are annoying. More has been written on this elsewhere.
- All such quantity types just mentioned boil down to `EMACS_INT`, which is 32 bits on 32-bit machines and 64 bits on 64-bit machines. This is guaranteed to be the same size as Lisp objects of type 'int', and (as far as I can tell) of `size_t` (unsigned!) and `ssize_t`. The only type below that is not an `EMACS_INT` is `Hashcode`, which is an unsigned value of the same size as `EMACS_INT`.
- Type names should be relatively short (no more than 10 characters or so), with the first letter capitalized and no underscores if they can at all be avoided.
- "count" == a zero-based measurement of some quantity. Includes sizes, offsets, and indexes.
- "bpos" == a one-based measurement of a position in a buffer. "Charbpos" and "Bytebpos" count text in the buffer, rather than bytes in memory; thus `Bytebpos` does not directly correspond to the memory representation. Use "Memnpos" for this.
- "Char" refers to internal-format characters, not to the C type "char", which is really a byte.

For the actual name changes, see the script below.

I ran the following script to do the conversion. (NOTE: This script is idempotent. You can safely run it multiple times and it will not screw up previous results – in fact, it will do nothing if nothing has changed. Thus, it can be run repeatedly as necessary to handle patches coming in from old workspaces, or old branches.) There are two tags, just before and

just after the change: ‘pre-integral-type-rename’ and ‘post-integral-type-rename’. When merging code from the main trunk into a branch, the best thing to do is first merge up to ‘pre-integral-type-rename’, then apply the script and associated changes, then merge from ‘post-integral-type-change’ to the present. (Alternatively, just do the merging in one operation; but you may then have a lot of conflicts needing to be resolved by hand.)

Script ‘fixtypes.sh’ follows:

```

----- cut -----
files="*. [ch] s/*.h m/*.h config.h.in ../configure.in Makefile.in.in ../lib-src/*. [ch]
gr Memory_Count Bytecount $files
gr Lstream_Data_Count Bytecount $files
gr Element_Count Elemcount $files
gr Hash_Code Hashcode $files
gr extcount bytecount $files
gr bufpos charbpos $files
gr bytind bytebpos $files
gr memind membpos $files
gr bufbyte intbyte $files
gr Extcount Bytecount $files
gr Bufpos Charbpos $files
gr Bytind Bytebpos $files
gr Memind Membpos $files
gr Bufbyte Intbyte $files
gr EXTCOUNT BYTECOUNT $files
gr BUFPOS CHARBPOS $files
gr BYTIND BYTEBPOS $files
gr MEMIND MEMBPOS $files
gr BUFBYTE INTBYTE $files
gr MEMORY_COUNT BYTECOUNT $files
gr LSTREAM_DATA_COUNT BYTECOUNT $files
gr ELEMENT_COUNT ELEM COUNT $files
gr HASH_CODE HASHCODE $files
----- cut -----

```

The ‘gr’ script, and the scripts it uses, are documented in ‘README.global-renaming’, because if placed in this file they would need to have their @ characters doubled, meaning you couldn’t easily cut and paste from the source.

In addition to those programs, I needed to fix up a few other things, particularly relating to the duplicate definitions of types, now that some types merged with others. Specifically:

1. in lisp.h, removed duplicate declarations of Bytecount. The changed code should now look like this: (In each code snippet below, the first and last lines are the same as the original, as are all lines outside of those lines. That allows you to locate the section to be replaced, and replace the stuff in that section, verifying that there isn’t anything new added that would need to be kept.)

```

----- snip -----
/* Counts of bytes or chars */
typedef EMACS_INT Bytecount;

```

```

typedef EMACS_INT Charcount;

/* Counts of elements */
typedef EMACS_INT Elemcount;

/* Hash codes */
typedef unsigned long Hashcode;

/* ----- dynamic arrays ----- */
----- snip -----

```

2. in `lstream.h`, removed duplicate declaration of `Bytecount`. Rewrote the comment about this type. The changed code should now look like this:

```

----- snip -----
#endif

/* The have been some arguments over the what the type should be that
   specifies a count of bytes in a data block to be written out or read in,
   using Lstream_read(), Lstream_write(), and related functions.
   Originally it was long, which worked fine; Martin "corrected" these to
   size_t and ssize_t on the grounds that this is theoretically cleaner and
   is in keeping with the C standards. Unfortunately, this practice is
   horribly error-prone due to design flaws in the way that mixed
   signed/unsigned arithmetic happens. In fact, by doing this change,
   Martin introduced a subtle but fatal error that caused the operation of
   sending large mail messages to the SMTP server under Windows to fail.
   By putting all values back to be signed, avoiding any signed/unsigned
   mixing, the bug immediately went away. The type then in use was
   Lstream_Data_Count, so that it be reverted cleanly if a vote came to
   that. Now it is Bytecount.
```

Some earlier comments about why the type must be signed: This MUST BE SIGNED, since it also is used in functions that return the number of bytes actually read to or written from in an operation, and these functions can return -1 to signal error.

Note that the standard Unix `read()` and `write()` functions define the count going in as a `size_t`, which is UNSIGNED, and the count going out as an `ssize_t`, which is SIGNED. This is a horrible design flaw. Not only is it highly likely to lead to logic errors when a -1 gets interpreted as a large positive number, but operations are bound to fail in all sorts of horrible ways when a number in the upper-half of the `size_t` range is passed in -- this number is unrepresentable as an `ssize_t`, so code that checks to see how many bytes are actually written (which is mandatory if you are dealing with certain types of devices) will get completely screwed up.

```

    --ben
*/

```

```

typedef enum lstream_buffering

```

```

----- snip -----

```

3. in `dumper.c`, there are four places, all inside of `switch()` statements, where `XD_BYTECOUNT` appears twice as a case tag. In each case, the two case blocks contain identical code, and you should *\*REMOVE THE SECOND\** and leave the first.

## 9.2 Text/Char Type Renaming

The purpose of this was

1. To distinguish between “charptr” when it refers to operations on the pointer itself and when it refers to operations on text
2. To use consistent naming for everything referring to internal format, i.e.

```

Itext == text in internal format

```

```

Ibyte == a byte in such text

```

```

Ichar == a char as represented in internal character format

```

Thus e.g.

```

set_charptr_emchar -> set_itext_ichar

```

This was done using a script like this:

```

files="*. [ch] s/*.h m/*.h config.h.in ../configure.in Makefile.in.in ../lib-src/*. [ch]
gr Intbyte Ibyte $files
gr INTBYTE IBYTE $files
gr intbyte ibyte $files
gr EMCHAR ICHAR $files
gr emchar ichar $files
gr Emchar Ichar $files
gr INC_CHARPTR INC_IBYTEPTR $files
gr DEC_CHARPTR DEC_IBYTEPTR $files
gr VALIDATE_CHARPTR VALIDATE_IBYTEPTR $files
gr valid_charptr valid_ibyteptr $files
gr CHARPTR ITEXT $files
gr charptr itext $files
gr Charptr Itext $files

```

See above for the source to ‘gr’.

As in the `integral-types` change, there are pre and post tags before and after the change:

```

pre-internal-format-textual-renaming

```

```

post-internal-format-textual-renaming

```

When merging a large branch, follow the same sort of procedure documented above, using these tags – essentially sync up to the pre tag, then apply the script yourself, then sync from the post tag to the present. You can probably do the same if you don’t have a separate workspace, but do have lots of outstanding changes and you’d rather not just merge all the textual changes directly. Use something like this:

(WARNING: I'm not a CVS guru; before trying this, or any large operation that might potentially mess things up, \*DEFINITELY\* make a backup of your existing workspace.)

```
cup -r pre-internal-format-textual-renaming
<apply script>
cup -A -j post-internal-format-textual-renaming -j HEAD
```

This might also work:

```
cup -j pre-internal-format-textual-renaming
<apply script>
cup -j post-internal-format-textual-renaming -j HEAD
```

ben

The following is a script to go in the opposite direction:

```
files="*. [ch] s/*.h m/*.h config.h.in ../configure.in Makefile.in.in ../lib-src/*. [ch]
```

```
# Evidently Perl considers _ to be a word char ala \b, even though XEmacs
# doesn't. We need to be careful here with ibyte/ichar because of words
# like Richard, eicharlen(), multibyte, HIBYTE, etc.
```

```
gr Ibyte Intbyte $files
gr '\bIBYTE' INTBYTE $files
gr '\bibyte' intbyte $files
gr '\bICHR' EMCHAR $files
gr '\bichar' emchar $files
gr '\bIchr' Emchar $files
gr '\bIBYTEPTR' CHARPTR $files
gr '\bibyteptr' charptr $files
gr '\bITEXT' CHARPTR $files
gr '\bitext' charptr $files
gr '\bItext' CHARPTR $files
```

```
gr '_IBYTE' _INTBYTE $files
gr '_ibyte' _intbyte $files
gr '_ICHR' _EMCHAR $files
gr '_ichar' _emchar $files
gr '_Ichr' _Emchar $files
gr '_IBYTEPTR' _CHARPTR $files
gr '_ibyteptr' _charptr $files
gr '_ITEXT' _CHARPTR $files
gr '_itext' _charptr $files
gr '_Itext' _CHARPTR $files
```



## 10 Rules When Writing New C Code

The XEmacs C Code is extremely complex and intricate, and there are many rules that are more or less consistently followed throughout the code. Many of these rules are not obvious, so they are explained here. It is of the utmost importance that you follow them. If you don't, you may get something that appears to work, but which will crash in odd situations, often in code far away from where the actual breakage is.

### 10.1 A Reader's Guide to XEmacs Coding Conventions

Of course the low-level implementation language of XEmacs is C, but much of that uses the Lisp engine to do its work. However, because the code is “inside” of the protective containment shell around the “reactor core,” you'll see lots of complex “plumbing” needed to do the work and “safety mechanisms,” whose failure results in a meltdown. This section provides a quick overview (or review) of the various components of the implementation of Lisp objects.

Two typographic conventions help to identify C objects that implement Lisp objects. The first is that capitalized identifiers, especially beginning with the letters ‘Q’, ‘V’, ‘F’, and ‘S’, for C variables and functions, and C macros with beginning with the letter ‘X’, are used to implement Lisp. The second is that where Lisp uses the hyphen ‘-’ in symbol names, the corresponding C identifiers use the underscore ‘\_’. Of course, since XEmacs Lisp contains interfaces to many external libraries, those external names will follow the coding conventions their authors chose, and may overlap the “XEmacs name space.” However these cases are usually pretty obvious.

All Lisp objects are handled indirectly. The `Lisp_Object` type is usually a pointer to a structure, except for a very small number of types with immediate representations (currently characters and integers). However, these types cannot be directly operated on in C code, either, so they can also be considered indirect. Types that do not have an immediate representation always have a C typedef `Lisp_type` for a corresponding structure.

In older code, it was common practice to pass around pointers to `Lisp_type`, but this is now deprecated in favor of using `Lisp_Object` for all function arguments and return values that are Lisp objects. The `Xtype` macro is used to extract the pointer and cast it to `(Lisp_type *)` for the desired type.

**Convention:** macros whose names begin with ‘X’ operate on `Lisp_Objects` and do no type-checking. Many such macros are type extractors, but others implement Lisp operations in C (*e.g.*, `XCAR` implements the Lisp `car` function). These are unsafe, and must only be used where types of all data have already been checked. Such macros are only applied to `Lisp_Objects`. In internal implementations where the pointer has already been converted, the structure is operated on directly using the C `->` member access operator.

The `type P`, `CHECK_type`, and `CONCHECK_type` macros are used to test types. The first returns a Boolean value, and the latter signal errors. (The ‘`CONCHECK`’ variety allows execution to be `CONTinued` under some circumstances, thus the name.) Functions which expect to be passed user data invariably call ‘`CHECK`’ macros on arguments.

There are many types of specialized Lisp objects implemented in C, but the most pervasive type is the **symbol**. Symbols are used as identifiers, variables, and functions.

**Convention:** Global variables whose names begin with ‘Q’ are constants whose value is a symbol. The name of the variable should be derived from the name of the symbol using the same rules as for Lisp primitives. Such variables allow the C code to check whether a particular `Lisp_Object` is equal to a given symbol. Symbols are Lisp objects, so these variables may be passed to Lisp primitives. (An alternative to the use of ‘Q...’ variables is to call the `intern` function at initialization in the `vars_of_module` function, which is hardly less efficient.)

**Convention:** Global variables whose names begin with ‘V’ are variables that contain Lisp objects. The convention here is that all global variables of type `Lisp_Object` begin with ‘V’, and no others do (not even integer and boolean variables that have Lisp equivalents). Most of the time, these variables have equivalents in Lisp, which are defined via the ‘DEFVAR’ family of macros, but some don’t. Since the variable’s value is a `Lisp_Object`, it can be passed to Lisp primitives.

The implementation of Lisp primitives is more complex. **Convention:** Global variables with names beginning with ‘S’ contain a structure that allows the Lisp engine to identify and call a C function. In modern versions of XEmacs, these identifiers are almost always completely hidden in the `DEFUN` and `SUBR` macros, but you will encounter them if you look at very old versions of XEmacs or at GNU Emacs. **Convention:** Functions with names beginning with ‘F’ implement Lisp primitives. Of course all their arguments and their return values must be Lisp-Objects. (This is hidden in the `DEFUN` macro.)

## 10.2 General Coding Rules

The C code is actually written in a dialect of C called **Clean C**, meaning that it can be compiled, mostly warning-free, with either a C or C++ compiler. Coding in Clean C has several advantages over plain C. C++ compilers are more nit-picking, and a number of coding errors have been found by compiling with C++. The ability to use both C and C++ tools means that a greater variety of development tools are available to the developer.

Every module includes ‘<config.h>’ (angle brackets so that ‘--srcdir’ works correctly; ‘config.h’ may or may not be in the same directory as the C sources) and ‘lisp.h’. ‘config.h’ must always be included before any other header files (including system header files) to ensure that certain tricks played by various ‘s/’ and ‘m/’ files work out correctly.

When including header files, always use angle brackets, not double quotes, except when the file to be included is always in the same directory as the including file. If either file is a generated file, then that is not likely to be the case. In order to understand why we have this rule, imagine what happens when you do a build in the source directory using ‘./configure’ and another build in another directory using ‘./work/configure’. There will be two different ‘config.h’ files. Which one will be used if you ‘#include "config.h"’?

Almost every module contains a `syms_of_*`() function and a `vars_of_*`() function. The former declares any Lisp primitives you have defined and defines any symbols you will be using. The latter declares any global Lisp variables you have added and initializes global C variables in the module. **Important:** There are stringent requirements on exactly what can go into these functions. See the comment in ‘emacs.c’. The reason for this is to avoid obscure unwanted interactions during initialization. If you don’t follow these rules, you’ll be sorry! If you want to do anything that isn’t allowed, create a `complex_vars_of_*`()

function for it. Doing this is tricky, though: you have to make sure your function is called at the right time so that all the initialization dependencies work out.

Declare each function of these kinds in `'symsinit.h'`. Make sure it's called in the appropriate place in `'emacs.c'`. You never need to include `'symsinit.h'` directly, because it is included by `'lisp.h'`.

**All global and static variables that are to be modified must be declared uninitialized.** This means that you may not use the “declare with initializer” form for these variables, such as `int some_variable = 0;`. The reason for this has to do with some kludges done during the dumping process: If possible, the initialized data segment is re-mapped so that it becomes part of the (unmodifiable) code segment in the dumped executable. This allows this memory to be shared among multiple running XEmacs processes. XEmacs is careful to place as much constant data as possible into initialized variables during the `'temacs'` phase.

**Please note:** This kludge only works on a few systems nowadays, and is rapidly becoming irrelevant because most modern operating systems provide **copy-on-write** semantics. All data is initially shared between processes, and a private copy is automatically made (on a page-by-page basis) when a process first attempts to write to a page of memory.

Formerly, there was a requirement that static variables not be declared inside of functions. This had to do with another hack along the same vein as what was just described: old USG systems put statically-declared variables in the initialized data space, so those header files had a `#define static` declaration. (That way, the data-segment remapping described above could still work.) This fails badly on static variables inside of functions, which suddenly become automatic variables; therefore, you weren't supposed to have any of them. This awful kludge has been removed in XEmacs because

1. almost all of the systems that used this kludge ended up having to disable the data-segment remapping anyway;
2. the only systems that didn't were extremely outdated ones;
3. this hack completely messed up inline functions.

The C source code makes heavy use of C preprocessor macros. One popular macro style is:

```
#define FOO(var, value) do {           \
    Lisp_Object FOO_value = (value);  \
    ... /* compute using FOO_value */ \
    (var) = bar;                       \
} while (0)
```

The `do { ... } while (0)` is a standard trick to allow `FOO` to have statement semantics, so that it can safely be used within an `if` statement in C, for example. Multiple evaluation is prevented by copying a supplied argument into a local variable, so that `FOO(var, fun(1))` only calls `fun` once.

Lisp lists are popular data structures in the C code as well as in Elisp. There are two sets of macros that iterate over lists. `EXTERNAL_LIST_LOOP_n` should be used when the list has been supplied by the user, and cannot be trusted to be acyclic and `nil`-terminated. A `malformed-list` or `circular-list` error will be generated if the list being iterated over is not entirely kosher. `LIST_LOOP_n`, on the other hand, is faster and less safe, and can be used only on trusted lists.

Related macros are `GET_EXTERNAL_LIST_LENGTH` and `GET_LIST_LENGTH`, which calculate the length of a list, and in the case of `GET_EXTERNAL_LIST_LENGTH`, validating the properness of the list. The macros `EXTERNAL_LIST_LOOP_DELETE_IF` and `LIST_LOOP_DELETE_IF` delete elements from a lisp list satisfying some predicate.

### 10.3 Writing Lisp Primitives

Lisp primitives are Lisp functions implemented in C. The details of interfacing the C function so that Lisp can call it are handled by a few C macros. The only way to really understand how to write new C code is to read the source, but we can explain some things here.

An example of a special form is the definition of `prog1`, from `'eval.c'`. (An ordinary function would have the same general appearance.)

```
DEFUN ("prog1", Fprog1, 1, UNEVALLED, 0, /*
  Similar to `progn', but the value of the first form is returned.
  \(\prog1 FIRST BODY...): All the arguments are evaluated sequentially.
  The value of FIRST is saved during evaluation of the remaining args,
  whose values are discarded.
  */)
  (args)
{
  /* This function can GC */
  REGISTER Lisp_Object val, form, tail;
  struct gcpro gcpro1;

  val = Feval (XCAR (args));

  GCPRO1 (val);

  LIST_LOOP_3 (form, XCDR (args), tail)
    Feval (form);

  UNGCPRO;
  return val;
}
```

Let's start with a precise explanation of the arguments to the `DEFUN` macro. Here is a template for them:

```
DEFUN (lname, fname, min_args, max_args, interactive, /*
  docstring
  */)
  (arglist )
```

**lname** This string is the name of the Lisp symbol to define as the function name; in the example above, it is `"prog1"`.

**fname** This is the C function name for this function. This is the name that is used in C code for calling the function. The name is, by convention, 'F' prepended to the Lisp name, with all dashes ('-') in the Lisp name changed to underscores. Thus, to call this function from C code, call `Fprog1`. Remember that the arguments are of type `Lisp_Object`; various macros and functions for creating values of type `Lisp_Object` are declared in the file `'lisp.h'`.

Primitives whose names are special characters (e.g. `+` or `<`) are named by spelling out, in some fashion, the special character: e.g. `Fplus()` or `Fless()`.

Primitives whose names begin with normal alphanumeric characters but also contain special characters are spelled out in some creative way, e.g. `let*` becomes `FletX()`.

Each function also has an associated structure that holds the data for the subr object that represents the function in Lisp. This structure conveys the Lisp symbol name to the initialization routine that will create the symbol and store the subr object as its definition. The C variable name of this structure is always 'S' prepended to the `fname`. You hardly ever need to be aware of the existence of this structure, since `DEFUN` plus `DEFSUBR` takes care of all the details.

**min\_args** This is the minimum number of arguments that the function requires. The function `prog1` allows a minimum of one argument.

**max\_args** This is the maximum number of arguments that the function accepts, if there is a fixed maximum. Alternatively, it can be `UNEVALLED`, indicating a special form that receives unevaluated arguments, or `MANY`, indicating an unlimited number of evaluated arguments (the C equivalent of `&rest`). Both `UNEVALLED` and `MANY` are macros. If `max_args` is a number, it may not be less than `min_args` and it may not be greater than 8. (If you need to add a function with more than 8 arguments, use the `MANY` form. Resist the urge to edit the definition of `DEFUN` in 'lisp.h'. If you do it anyways, make sure to also add another clause to the switch statement in `primitive_funcall()`.)

**interactive**

This is an interactive specification, a string such as might be used as the argument of `interactive` in a Lisp function. In the case of `prog1`, it is 0 (a null pointer), indicating that `prog1` cannot be called interactively. A value of "" indicates a function that should receive no arguments when called interactively.

**docstring**

This is the documentation string. It is written just like a documentation string for a function defined in Lisp; in particular, the first line should be a single sentence. Note how the documentation string is enclosed in a comment, none of the documentation is placed on the same lines as the comment-start and comment-end characters, and the comment-start characters are on the same line as the interactive specification. 'make-docfile', which scans the C files for documentation strings, is very particular about what it looks for, and will not properly extract the doc string if it's not in this exact format.

In order to make both 'etags' and 'make-docfile' happy, make sure that the `DEFUN` line contains the `lname` and `fname`, and that the comment-start characters for the doc string are on the same line as the interactive specification, and put a newline directly after them (and before the comment-end characters).

**arglist**

This is the comma-separated list of arguments to the C function. For a function with a fixed maximum number of arguments, provide a C argument for each Lisp argument. In this case, unlike regular C functions, the types of the arguments are not declared; they are simply always of type `Lisp_Object`.

The names of the C arguments will be used as the names of the arguments to the Lisp primitive as displayed in its documentation, modulo the same concerns described above for `F...` names (in particular, underscores in the C arguments become dashes in the Lisp arguments).

There is one additional kludge: A trailing ‘\_’ on the C argument is discarded when forming the Lisp argument. This allows C language reserved words (like `default`) or global symbols (like `dirname`) to be used as argument names without compiler warnings or errors.

A Lisp function with `max_args = UNEVALLED` is a **special form**; its arguments are not evaluated. Instead it receives one argument of type `Lisp_Object`, a (Lisp) list of the unevaluated arguments, conventionally named (`args`).

When a Lisp function has no upper limit on the number of arguments, specify `max_args = MANY`. In this case its implementation in C actually receives exactly two arguments: the number of Lisp arguments (an `int`) and the address of a block containing their values (a `Lisp_Object *`). In this case only are the C types specified in the `arglist`: (`int nargs, Lisp_Object *args`).

Within the function `Fprog1` itself, note the use of the macros `GCPR01` and `UNGCPRO`. `GCPR01` is used to “protect” a variable from garbage collection—to inform the garbage collector that it must look in that variable and regard the object pointed at by its contents as an accessible object. This is necessary whenever you call `Feval` or anything that can directly or indirectly call `Feval` (this includes the `QUIT` macro!). At such a time, any Lisp object that you intend to refer to again must be protected somehow. `UNGCPRO` cancels the protection of the variables that are protected in the current function. It is necessary to do this explicitly.

The macro `GCPR01` protects just one local variable. If you want to protect two, use `GCPR02` instead; repeating `GCPR01` will not work. Macros `GCPR03` and `GCPR04` also exist.

These macros implicitly use local variables such as `gcpro1`; you must declare these explicitly, with type `struct gcpro`. Thus, if you use `GCPR02`, you must declare `gcpro1` and `gcpro2`.

Note also that the general rule is **caller-protects**; i.e. you are only responsible for protecting those Lisp objects that you create. Any objects passed to you as arguments should have been protected by whoever created them, so you don’t in general have to protect them.

In particular, the arguments to any Lisp primitive are always automatically `GCPR0ed`, when called “normally” from Lisp code or bytecode. So only a few Lisp primitives that are called frequently from C code, such as `Fprogn` protect their arguments as a service to their caller. You don’t need to protect your arguments when writing a new `DEFUN`.

`GCPR0ing` is perhaps the trickiest and most error-prone part of XEmacs coding. It is **extremely** important that you get this right and use a great deal of discipline when writing this code. See [Section 14.3 \[GCPR0ing\]](#), page 86, for full details on how to do this.

What `DEFUN` actually does is declare a global structure of type `Lisp_Subr` whose name begins with capital ‘SF’ and which contains information about the primitive (e.g. a pointer to the function, its minimum and maximum allowed arguments, a string describing its Lisp name); `DEFUN` then begins a normal C function declaration using the `F...` name. The Lisp subr object that is the function definition of a primitive (i.e. the object in the function slot of the symbol that names the primitive) actually points to this ‘SF’ structure; when `Feval` encounters a subr, it looks in the structure to find out how to call the C function.

Defining the C function is not enough to make a Lisp primitive available; you must also create the Lisp symbol for the primitive (the symbol is **interned**; see [Section 18.2 \[Obarrays\]](#),

page 121) and store a suitable `subr` object in its function cell. (If you don't do this, the primitive won't be seen by Lisp code.) The code looks like this:

```
DEFSUBR (fname);
```

Here `fname` is the same name you used as the second argument to `DEFUN`.

This call to `DEFSUBR` should go in the `syms_of_*`(`*`) function at the end of the module. If no such function exists, create it and make sure to also declare it in `'symsinit.h'` and call it from the appropriate spot in `main()`. See [Section 10.2 \[General Coding Rules\]](#), page 38.

Note that C code cannot call functions by name unless they are defined in C. The way to call a function written in Lisp from C is to use `Ffuncall`, which embodies the Lisp function `funcall`. Since the Lisp function `funcall` accepts an unlimited number of arguments, in C it takes two: the number of Lisp-level arguments, and a one-dimensional array containing their values. The first Lisp-level argument is the Lisp function to call, and the rest are the arguments to pass to it. Since `Ffuncall` can call the evaluator, you must protect pointers from garbage collection around the call to `Ffuncall`. (However, `Ffuncall` explicitly protects all of its parameters, so you don't have to protect any pointers passed as parameters to it.)

The C functions `call0`, `call1`, `call2`, and so on, provide handy ways to call a Lisp function conveniently with a fixed number of arguments. They work by calling `Ffuncall`.

`'eval.c'` is a very good file to look through for examples; `'lisp.h'` contains the definitions for important macros and functions.

## 10.4 Writing Good Comments

Comments are a lifeline for programmers trying to understand tricky code. In general, the less obvious it is what you are doing, the more you need a comment, and the more detailed it needs to be. You should always be on guard when you're writing code for stuff that's tricky, and should constantly be putting yourself in someone else's shoes and asking if that person could figure out without much difficulty what's going on. (Assume they are a competent programmer who understands the essentials of how the XEmacs code is structured but doesn't know much about the module you're working on or any algorithms you're using.) If you're not sure whether they would be able to, add a comment. Always err on the side of more comments, rather than less.

Generally, when making comments, there is no need to attribute them with your name or initials. This especially goes for small, easy-to-understand, non-opinionated ones. Also, comments indicating where, when, and by whom a file was changed are *strongly* discouraged, and in general will be removed as they are discovered. This is exactly what `'ChangeLogs'` are there for. However, it can occasionally be useful to mark exactly where (but not when or by whom) changes are made, particularly when making small changes to a file imported from elsewhere. These marks help when later on a newer version of the file is imported and the changes need to be merged. (If everything were always kept in CVS, there would be no need for this. But in practice, this often doesn't happen, or the CVS repository is later on lost or unavailable to the person doing the update.)

When putting in an explicit opinion in a comment, you should *always* attribute it with your name, and optionally the date. This also goes for long, complex comments explaining in detail the workings of something – by putting your name there, you make it possible

for someone who has questions about how that thing works to determine who wrote the comment so they can write to them. Preferably, use your actual name and not your initials, unless your initials are generally recognized (e.g. ‘jwz’). You can use only your first name if it’s obvious who you are; otherwise, give first and last name. If you’re not a regular contributor, you might consider putting your email address in – it may be in the ChangeLog, but after awhile ChangeLogs have a tendency of disappearing or getting muddled. (E.g. your comment may get copied somewhere else or even into another program, and tracking down the proper ChangeLog may be very difficult.)

If you come across an opinion that is not or no longer valid, or you come across any comment that no longer applies but you want to keep it around, enclose it in ‘[[ ’ and ‘ ]]’ marks and add a comment afterwards explaining why the preceding comment is no longer valid. Put your name on this comment, as explained above.

Just as comments are a lifeline to programmers, incorrect comments are death. If you come across an incorrect comment, **immediately** correct it or flag it as incorrect, as described in the previous paragraph. Whenever you work on a section of code, *always* make sure to update any comments to be correct – or, at the very least, flag them as incorrect.

To indicate a "todo" or other problem, use four pound signs – i.e. ‘####’.

## 10.5 Adding Global Lisp Variables

Global variables whose names begin with ‘Q’ are constants whose value is a symbol of a particular name. The name of the variable should be derived from the name of the symbol using the same rules as for Lisp primitives. These variables are initialized using a call to `defsymbol()` in the `syms_of_*` function. (This call interns a symbol, sets the C variable to the resulting Lisp object, and calls `staticpro()` on the C variable to tell the garbage-collection mechanism about this variable. What `staticpro()` does is add a pointer to the variable to a large global array; when garbage-collection happens, all pointers listed in the array are used as starting points for marking Lisp objects. This is important because it’s quite possible that the only current reference to the object is the C variable. In the case of symbols, the `staticpro()` doesn’t matter all that much because the symbol is contained in `obarray`, which is itself `staticpro()`ed. However, it’s possible that a naughty user could do something like uninterning the symbol out of `obarray` or even setting `obarray` to a different value [although this is likely to make XEmacs crash!].)

**Please note:** It is potentially deadly if you declare a ‘Q...’ variable in two different modules. The two calls to `defsymbol()` are no problem, but some linkers will complain about multiply-defined symbols. The most insidious aspect of this is that often the link will succeed anyway, but then the resulting executable will sometimes crash in obscure ways during certain operations!

To avoid this problem, declare any symbols with common names (such as `text`) that are not obviously associated with this particular module in the file `general-slots.h`. The “-slots” suffix indicates that this is a file that is included multiple times in `general.c`. Redefinition of preprocessor macros allows the effects to be different in each context, so this is actually more convenient and less error-prone than doing it in your module.

Global variables whose names begin with ‘V’ are variables that contain Lisp objects. The convention here is that all global variables of type `Lisp_Object` begin with ‘V’, and all others don’t (including integer and boolean variables that have Lisp equivalents). Most

of the time, these variables have equivalents in Lisp, but some don't. Those that do are declared this way by a call to `DEFVAR_LISP()` in the `vars_of_*` initializer for the module. What this does is create a special **symbol-value-forward** Lisp object that contains a pointer to the C variable, intern a symbol whose name is as specified in the call to `DEFVAR_LISP()`, and set its value to the symbol-value-forward Lisp object; it also calls `staticpro()` on the C variable to tell the garbage-collection mechanism about the variable. When `eval` (or actually `symbol-value`) encounters this special object in the process of retrieving a variable's value, it follows the indirection to the C variable and gets its value. `setq` does similar things so that the C variable gets changed.

Whether or not you `DEFVAR_LISP()` a variable, you need to initialize it in the `vars_of_*` function; otherwise it will end up as all zeroes, which is the integer 0 (*not nil*), and this is probably not what you want. Also, if the variable is not `DEFVAR_LISP()`ed, you must call `staticpro()` on the C variable in the `vars_of_*` function. Otherwise, the garbage-collection mechanism won't know that the object in this variable is in use, and will happily collect it and reuse its storage for another Lisp object, and you will be the one who's unhappy when you can't figure out how your variable got overwritten.

## 10.6 Proper Use of Unsigned Types

Avoid using `unsigned int` and `unsigned long` whenever possible. Unsigned types are viral – any arithmetic or comparisons involving mixed signed and unsigned types are automatically converted to unsigned, which is almost certainly not what you want. Many subtle and hard-to-find bugs are created by careless use of unsigned types. In general, you should almost *never* use an unsigned type to hold a regular quantity of any sort. The only exceptions are

1. When there's a reasonable possibility you will actually need all 32 or 64 bits to store the quantity.
2. When calling existing API's that require unsigned types. In this case, you should still do all manipulation using signed types, and do the conversion at the very threshold of the API call.
3. In existing code that you don't want to modify because you don't maintain it.
4. In bit-field structures.

Other reasonable uses of `unsigned int` and `unsigned long` are representing non-quantities – e.g. bit-oriented flags and such.

## 10.7 Coding for Mule

Although Mule support is not compiled by default in XEmacs, many people are using it, and we consider it crucial that new code works correctly with multibyte characters. This is not hard; it is only a matter of following several simple user-interface guidelines. Even if you never compile with Mule, with a little practice you will find it quite easy to code Mule-correctly.

Note that these guidelines are not necessarily tied to the current Mule implementation; they are also a good idea to follow on the grounds of code generalization for future I18N work.

### 10.7.1 Character-Related Data Types

First, let's review the basic character-related datatypes used by XEmacs. Note that some of the separate `typedefs` are not mandatory, but they improve clarity of code a great deal, because one glance at the declaration can tell the intended use of the variable.

**Ichar** An **Ichar** holds a single Emacs character.

Obviously, the equality between characters and bytes is lost in the Mule world. Characters can be represented by one or more bytes in the buffer, and **Ichar** is a C type large enough to hold any character.

Without Mule support, an **Ichar** is equivalent to an `unsigned char`.

**Ibyte** The data representing the text in a buffer or string is logically a set of **Ibytes**. XEmacs does not work with the same character formats all the time; when reading characters from the outside, it decodes them to an internal format, and likewise encodes them when writing. **Ibyte** (in fact `unsigned char`) is the basic unit of XEmacs internal buffers and strings format. An **Ibyte \*** is the type that points at text encoded in the variable-width internal encoding.

One character can correspond to one or more **Ibytes**. In the current Mule implementation, an ASCII character is represented by the same **Ibyte**, and other characters are represented by a sequence of two or more **Ibytes**.

Without Mule support, there are exactly 256 characters, implicitly Latin-1, and each character is represented using one **Ibyte**, and there is a one-to-one correspondence between **Ibytes** and **Ichars**.

**Charxpos**

**Charbpos**

**Charcount**

A **Charbpos** represents a character position in a buffer. A **Charcount** represents a number (count) of characters. Logically, subtracting two **Charbpos** values yields a **Charcount** value. When representing a character position in a string, we just use **Charcount** directly. The reason for having a separate typedef for buffer positions is that they are 1-based, whereas string positions are 0-based and hence string counts and positions can be freely intermixed (a string position is equivalent to the count of characters from the beginning). When representing a character position that could be either in a buffer or string (for example, in the extent code), **Charxpos** is used. Although all of these are `typedefed` to `EMACS_INT`, we use them in preference to `EMACS_INT` to make it clear what sort of position is being used.

**Charxpos**, **Charbpos** and **Charcount** values are the only ones that are ever visible to Lisp.

**Bytexpos**

**Bytecount**

A **Bytebpos** represents a byte position in a buffer. A **Bytecount** represents the distance between two positions, in bytes. Byte positions in strings use **Bytecount**, and for byte positions that can be either in a buffer or string, **Bytexpos** is used. The relationship between **Bytexpos**, **Bytebpos** and

`Bytecount` is the same as the relationship between `Charxpos`, `Charbpos` and `Charcount`.

**Extbyte** When dealing with the outside world, XEmacs works with `Extbytes`, which are equivalent to `char`. The distance between two `Extbytes` is a `Bytecount`, since external text is a byte-by-byte encoding. `Extbytes` occur mainly at the transition point between internal text and external functions. XEmacs code should not, if it can possibly avoid it, do any actual manipulation using external text, since its format is completely unpredictable (it might not even be ASCII-compatible).

### 10.7.2 Working With Character and Byte Positions

Now that we have defined the basic character-related types, we can look at the macros and functions designed for work with them and for conversion between them. Most of these macros are defined in `'buffer.h'`, and we don't discuss all of them here, but only the most important ones. Examining the existing code is the best way to learn about them.

#### `MAX_ICHAR_LEN`

This preprocessor constant is the maximum number of buffer bytes to represent an Emacs character in the variable width internal encoding. It is useful when allocating temporary strings to keep a known number of characters. For instance:

```
{
  Charcount cclen;
  ...
  {
    /* Allocate place for cclen characters. */
    Ibyte *buf = (Ibyte *) alloca (cclen * MAX_ICHAR_LEN);
    ...
  }
}
```

If you followed the previous section, you can guess that, logically, multiplying a `Charcount` value with `MAX_ICHAR_LEN` produces a `Bytecount` value.

In the current Mule implementation, `MAX_ICHAR_LEN` equals 4. Without Mule, it is 1.

#### `itext_ichar`

##### `set_itext_ichar`

The `itext_ichar` macro takes a `Ibyte` pointer and returns the `Ichar` stored at that position. If it were a function, its prototype would be:

```
Ichar itext_ichar (Ibyte *p);
```

`set_itext_ichar` stores an `Ichar` to the specified byte position. It returns the number of bytes stored:

```
Bytecount set_itext_ichar (Ibyte *p, Ichar c);
```

It is important to note that `set_itext_ichar` is safe only for appending a character at the end of a buffer, not for overwriting a character in the middle. This is because the width of characters varies, and `set_itext_ichar` cannot resize the string if it writes, say, a two-byte character where a single-byte character used to reside.

A typical use of `set_itext_ichar` can be demonstrated by this example, which copies characters from buffer `buf` to a temporary string of `Ibyte`s.

```
{
  Charbpos pos;
  for (pos = beg; pos < end; pos++)
  {
    Ichar c = BUF_FETCH_CHAR (buf, pos);
    p += set_itext_ichar (buf, c);
  }
}
```

Note how `set_itext_ichar` is used to store the `Ichar` and increment the counter, at the same time.

`INC_IBYTEPTR`

`DEC_IBYTEPTR`

These two macros increment and decrement an `Ibyte` pointer, respectively. They will adjust the pointer by the appropriate number of bytes according to the byte length of the character stored there. Both macros assume that the memory address is located at the beginning of a valid character.

Without Mule support, `INC_IBYTEPTR (p)` and `DEC_IBYTEPTR (p)` simply expand to `p++` and `p--`, respectively.

`bytecount_to_charcount`

Given a pointer to a text string and a length in bytes, return the equivalent length in characters.

```
Charcount bytecount_to_charcount (Ibyte *p, Bytecount bc);
```

`charcount_to_bytecount`

Given a pointer to a text string and a length in characters, return the equivalent length in bytes.

```
Bytecount charcount_to_bytecount (Ibyte *p, Charcount cc);
```

`itext_n_addr`

Return a pointer to the beginning of the character offset `cc` (in characters) from `p`.

```
Ibyte *itext_n_addr (Ibyte *p, Charcount cc);
```

### 10.7.3 Conversion to and from External Data

When an external function, such as a C library function, returns a `char` pointer, you should almost never treat it as `Ibyte`. This is because these returned strings may contain 8bit characters which can be misinterpreted by XEmacs, and cause a crash. Likewise, when exporting a piece of internal text to the outside world, you should always convert it to an appropriate external encoding, lest the internal stuff (such as the infamous `\201` characters) leak out.

The interface to conversion between the internal and external representations of text are the numerous conversion macros defined in `'buffer.h'`. There used to be a fixed set of external formats supported by these macros, but now any coding system can be used with

them. The coding system alias mechanism is used to create the following logical coding systems, which replace the fixed external formats. The (dontusethis-set-symbol-value-handler) mechanism was enhanced to make this possible (more work on that is needed).

Often useful coding systems:

- Qbinary** This is the simplest format and is what we use in the absence of a more appropriate format. This converts according to the `binary` coding system:
- a. On input, bytes 0–255 are converted into (implicitly Latin-1) characters 0–255. A non-Mule xemacs doesn't really know about different character sets and the fonts to display them, so the bytes can be treated as text in different 1-byte encodings by simply setting the appropriate fonts. So in a sense, non-Mule xemacs is a multi-lingual editor if, for example, different fonts are used to display text in different buffers, faces, or windows. The specifier mechanism gives the user complete control over this kind of behavior.
  - b. On output, characters 0–255 are converted into bytes 0–255 and other characters are converted into '~'.
- Qnative** Format used for the external Unix environment—`argv[]`, stuff from `getenv()`, stuff from the `/etc/passwd` file, etc. This is encoded according to the encoding specified by the current locale.
- Qfile\_name** Format used for filenames. This is normally the same as `Qnative`, but the two should be distinguished for clarity and possible future separation – and also because `Qfile_name` can be changed using either the `file-name-coding-system` or `pathname-coding-system` (now obsolete) variables.
- Qctext** Compound-text format. This is the standard X11 format used for data stored in properties, selections, and the like. This is an 8-bit no-lock-shift ISO2022 coding system. This is a real coding system, unlike `Qfile_name`, which is user-definable.
- Qmswindows\_tstr** Used for external data in all MS Windows functions that are declared to accept data of type `LPTSTR` or `LPCSTR`. This maps to either `Qmswindows_multibyte` (a locale-specific encoding, same as `Qnative`) or `Qmswindows_unicode`, depending on whether XEmacs is being run under Windows 9X or Windows NT/2000/XP.

Many other coding systems are provided by default.

There are two fundamental macros to convert between external and internal format, as well as various convenience macros to simplify the most common operations.

`TO_INTERNAL_FORMAT` converts external data to internal format, and `TO_EXTERNAL_FORMAT` converts the other way around. The arguments each of these receives are a source type, a source, a sink type, a sink, and a coding system (or a symbol naming a coding system).

A typical call looks like

```
TO_EXTERNAL_FORMAT (LISP_STRING, str, C_STRING_MALLOC, ptr, Qfile_name);
```

which means that the contents of the lisp string `str` are written to a malloc'ed memory area which will be pointed to by `ptr`, after the function returns. The conversion will be

done using the `file-name` coding system, which will be controlled by the user indirectly by setting or binding the variable `file-name-coding-system`.

Some sources and sinks require two C variables to specify. We use some preprocessor magic to allow different source and sink types, and even different numbers of arguments to specify different types of sources and sinks.

So we can have a call that looks like

```
TO_INTERNAL_FORMAT (DATA, (ptr, len),
                   MALLOC, (ptr, len),
                   coding_system);
```

The parenthesized argument pairs are required to make the preprocessor magic work.

Here are the different source and sink types:

```
DATA, (ptr, len),
    input data is a fixed buffer of size len at address ptr
ALLOCA, (ptr, len),
    output data is placed in an alloca()ed buffer of size len pointed to by ptr
MALLOC, (ptr, len),
    output data is in a malloc()ed buffer of size len pointed to by ptr
C_STRING_ALLOCA, ptr,
    equivalent to ALLOCA (ptr, len_ignored) on output.
C_STRING_MALLOC, ptr,
    equivalent to MALLOC (ptr, len_ignored) on output
C_STRING, ptr,
    equivalent to DATA, (ptr, strlen/wcslen (ptr)) on input
LISP_STRING, string,
    input or output is a Lisp-Object of type string
LISP_BUFFER, buffer,
    output is written to (point) in lisp buffer bu er
LISP_LSTREAM, lstream,
    input or output is a Lisp-Object of type lstream
LISP_OPAQUE, object,
    input or output is a Lisp-Object of type opaque
```

A source type of `C_STRING` or a sink type of `C_STRING_ALLOCA` or `C_STRING_MALLOC` is appropriate where the external API is not `'\0'`-byte-clean – i.e. it expects strings to be terminated with a null byte. For external API's that are in fact `'\0'`-byte-clean, we should of course not use these.

The sinks to be specified must be lvalues, unless they are the lisp object types `LISP_LSTREAM` or `LISP_BUFFER`.

There is no problem using the same lvalue for source and sink.

Garbage collection is inhibited during these conversion operations, so it is OK to pass in data from Lisp strings using `XSTRING_DATA`.

For the sink types `ALLOCA` and `C_STRING_ALLOCA`, the resulting text is stored in a stack-allocated buffer, which is automatically freed on returning from the function. However, the sink types `MALLOC` and `C_STRING_MALLOC` return `xmalloc()`ed memory. The caller is responsible for freeing this memory using `xfree()`.

Note that it doesn't make sense for `LISP_STRING` to be a source for `TO_INTERNAL_FORMAT` or a sink for `TO_EXTERNAL_FORMAT`. You'll get an assertion failure if you try.

99% of conversions involve raw data or Lisp strings as both source and sink, and usually data is output as `alloca()`, or sometimes `xmalloc()`. For this reason, convenience macros are defined for many types of conversions involving raw data and/or Lisp strings, especially when the output is an `alloca()`ed string. (When the destination is a Lisp string, there are other functions that should be used instead – `build_ext_string()` and `make_ext_string()`, for example.) The convenience macros are of two types – the older kind that store the result into a specified variable, and the newer kind that return the result. The newer kind of macros don't exist when the output is sized data, because that would have two return values. NOTE: All convenience macros are ultimately defined in terms of `TO_EXTERNAL_FORMAT` and `TO_INTERNAL_FORMAT`. Thus, any comments above about the workings of these macros also apply to all convenience macros.

A typical old-style convenience macro is

```
C_STRING_TO_EXTERNAL (in, out, codesys);
```

This is equivalent to

```
TO_EXTERNAL_FORMAT (C_STRING, in, C_STRING_ALLOCA, out, codesys);
```

but is easier to write and somewhat clearer, since it clearly identifies the arguments without the clutter of having the preprocessor types mixed in.

The new-style equivalent is `NEW_C_STRING_TO_EXTERNAL (src, codesys)`, which *returns* the converted data (still in `alloca()` space). This is far more convenient for most operations.

### 10.7.4 General Guidelines for Writing Mule-Aware Code

This section contains some general guidance on how to write Mule-aware code, as well as some pitfalls you should avoid.

*Never use `char` and `char *`.*

In XEmacs, the use of `char` and `char *` is almost always a mistake. If you want to manipulate an Emacs character from “C”, use `Ichar`. If you want to examine a specific octet in the internal format, use `Ibyte`. If you want a Lisp-visible character, use a `Lisp_Object` and `make_char`. If you want a pointer to move through the internal text, use `Ibyte *`. Also note that you almost certainly do not need `Ichar *`. Other typedefs to clarify the use of `char` are `Char_ASCII`, `Char_Binary`, `UChar_Binary`, and `CByte`.

*Be careful not to confuse `Charcount`, `Bytcount`, `Charbpos` and `Bytebpos`.*

The whole point of using different types is to avoid confusion about the use of certain variables. Lest this effect be nullified, you need to be careful about using the right types.

*Always convert external data*

It is extremely important to always convert external data, because XEmacs can crash if unexpected 8-bit sequences are copied to its internal buffers literally.

This means that when a system function, such as `readdir`, returns a string, you normally need to convert it using one of the conversion macros described in the previous chapter, before passing it further to Lisp.

Actually, most of the basic system functions that accept '\0'-terminated string arguments, like `stat()` and `open()`, have **encapsulated** equivalents that do the internal to external conversion themselves. The encapsulated equivalents have a `qxe_` prefix and have string arguments of type `Ibyte *`, and you can pass internally encoded data to them, often from a Lisp string using `XSTRING_DATA`. (A better design might be to provide versions that accept Lisp strings directly.)

Also note that many internal functions, such as `make_string`, accept `Ibytes`, which removes the need for them to convert the data they receive. This increases efficiency because that way external data needs to be decoded only once, when it is read. After that, it is passed around in internal format.

#### *Do all work in internal format*

External-formatted data is completely unpredictable in its format. It may be fixed-width Unicode (not even ASCII compatible); it may be a modal encoding, in which case some occurrences of (e.g.) the slash character may be part of two-byte Asian-language characters, and a naive attempt to split apart a pathname by slashes will fail; etc. Internal-format text should be converted to external format only at the point where an external API is actually called, and the first thing done after receiving external-format text from an external API should be to convert it to internal text.

### 10.7.5 An Example of Mule-Aware Code

As an example of Mule-aware code, we will analyze the `string` function, which conses up a Lisp string from the character arguments it receives. Here is the definition, pasted from `alloc.c`:

```
DEFUN ("string", Fstring, 0, MANY, 0, /*
Concatenate all the argument characters and make the result a string.
*/
      (int nargs, Lisp_Object *args))
{
  Ibyte *storage = alloca_array (Ibyte, nargs * MAX_ICHAR_LEN);
  Ibyte *p = storage;

  for (; nargs; nargs--, args++)
    {
      Lisp_Object lisp_char = *args;
      CHECK_CHAR_COERCE_INT (lisp_char);
      p += set_itext_ichar (p, XCHAR (lisp_char));
    }
  return make_string (storage, p - storage);
}
```

Now we can analyze the source line by line.

Obviously, `string` will be as long as there are arguments to the function. This is why we allocate `MAX_ICHAR_LEN * nargs` bytes on the stack, i.e. the worst-case number of bytes for `nargs` `Ichars` to fit in the string.

Then, the loop checks that each element is a character, converting integers in the process. Like many other functions in XEmacs, this function silently accepts integers where characters are expected, for historical and compatibility reasons. Unless you know what you are doing, `CHECK_CHAR` will also suffice. `XCHAR` (`lisp_char`) extracts the `Ichar` from the `Lisp_Object`, and `set_itext_ichar` stores it to storage, increasing `p` in the process.

Other instructive examples of correct coding under Mule can be found all over the XEmacs code. For starters, I recommend `Fnormalize_menu_item_name` in `menubar.c`. After you have understood this section of the manual and studied the examples, you can proceed writing new Mule-aware code.

### 10.7.6 Mule-izing Code

A lot of code is written without Mule in mind, and needs to be made Mule-correct or "Mule-ized". There is really no substitute for line-by-line analysis when doing this, but the following checklist can help:

- Check all uses of `XSTRING_DATA`.
- Check all uses of `build_string` and `make_string`.
- Check all uses of `tolower` and `toupper`.
- Check object print methods.
- Check for use of functions such as `write_c_string`, `write_fmt_string`, `stderr_out`, `stdout_out`.
- Check all occurrences of `char` and correct to one of the other typedefs described above.
- Check all existing uses of `TO_EXTERNAL_FORMAT`, `TO_INTERNAL_FORMAT`, and any convenience macros (grep for `'EXTERNAL_TO'`, `'TO_EXTERNAL'`, and `'TO_SIZED_EXTERNAL'`).
- In Windows code, string literals may need to be encapsulated with `XETEXT`.

## 10.8 Techniques for XEmacs Developers

To make a purified XEmacs, do: `make puremacs`. To make a quantified XEmacs, do: `make quantmacs`.

You simply can't dump Quantified and Purified images (unless using the portable dumper). Purify gets confused when `xemacs` frees memory in one process that was allocated in a *different* process on a different machine!. Run it like so:

```
temacs -batch -l loadup.el run-temacs xemacs-args...
```

Before you go through the trouble, are you compiling with all debugging and error-checking off? If not, try that first. Be warned that while Quantify is directly responsible for quite a few optimizations which have been made to XEmacs, doing a run which generates results which can be acted upon is not necessarily a trivial task.

Also, if you're still willing to do some runs make sure you configure with the `'--quantify'` flag. That will keep Quantify from starting to record data until after the loadup is completed and will shut off recording right before it shuts down (which generates enough bogus data

to throw most results off). It also enables three additional elisp commands: `quantify-start-recording-data`, `quantify-stop-recording-data` and `quantify-clear-data`.

If you want to make XEmacs faster, target your favorite slow benchmark, run a profiler like `Quantify`, `gprof`, or `tcov`, and figure out where the cycles are going. In many cases you can localize the problem (because a particular new feature or even a single patch elicited it). Don't hesitate to use brute force techniques like a global counter incremented at strategic places, especially in combination with other performance indications (*e.g.*, degree of buffer fragmentation into extents).

Specific projects:

- Make the garbage collector faster. Figure out how to write an incremental garbage collector.
- Write a compiler that takes bytecode and spits out C code. Unfortunately, you will then need a C compiler and a more fully developed module system.
- Speed up redisplay.
- Speed up syntax highlighting. It was suggested that “maybe moving some of the syntax highlighting capabilities into C would make a difference.” Wrong idea, I think. When processing one large file a particular low-level routine was being called 40 *million* times simply for *one* call to `newline-and-indent`. Syntax highlighting needs to be rewritten to use a reliable, fast parser, then to trust the pre-parsed structure, and only do re-highlighting locally to a text change. Modern machines are fast enough to implement such parsers in Lisp; but no machine will ever be fast enough to deal with quadratic (or worse) algorithms!
- Implement tail recursion in Emacs Lisp (hard!).

Unfortunately, Emacs Lisp is slow, and is going to stay slow. Function calls in elisp are especially expensive. Iterating over a long list is going to be 30 times faster implemented in C than in Elisp.

Heavily used small code fragments need to be fast. The traditional way to implement such code fragments in C is with macros. But macros in C are known to be broken.

Macro arguments that are repeatedly evaluated may suffer from repeated side effects or suboptimal performance.

Variable names used in macros may collide with caller's variables, causing (at least) unwanted compiler warnings.

In order to solve these problems, and maintain statement semantics, one should use the `do { ... } while (0)` trick while trying to reference macro arguments exactly once using local variables.

Let's take a look at this poor macro definition:

```
#define MARK_OBJECT(obj) \
    if (!marked_p (obj)) mark_object (obj), did_mark = 1
```

This macro evaluates its argument twice, and also fails if used like this:

```
if (flag) MARK_OBJECT (obj); else do_something();
```

A much better definition is

```
#define MARK_OBJECT(obj) do { \
    Lisp_Object mo_obj = (obj); \
```

```

    if (!marked_p (mo_obj))    \
    {                          \
        mark_object (mo_obj);  \
        did_mark = 1;         \
    }                          \
} while (0)

```

Notice the elimination of double evaluation by using the local variable with the obscure name. Writing safe and efficient macros requires great care. The one problem with macros that cannot be portably worked around is, since a C block has no value, a macro used as an expression rather than a statement cannot use the techniques just described to avoid multiple evaluation.

In most cases where a macro has function semantics, an inline function is a better implementation technique. Modern compiler optimizers tend to inline functions even if they have no `inline` keyword, and configure magic ensures that the `inline` keyword can be safely used as an additional compiler hint. Inline functions used in a single `.c` files are easy. The function must already be defined to be `static`. Just add another `inline` keyword to the definition.

```

inline static int
heavily_used_small_function (int arg)
{
    ...
}

```

Inline functions in header files are trickier, because we would like to make the following optimization if the function is *not* inlined (for example, because we're compiling for debugging). We would like the function to be defined externally exactly once, and each calling translation unit would create an external reference to the function, instead of including a definition of the inline function in the object code of every translation unit that uses it. This optimization is currently only available for `gcc`. But you don't have to worry about the trickiness; just define your inline functions in header files using this pattern:

```

INLINE_HEADER int
i_used_to_be_a_crufty_macro_but_look_at_me_now (int arg);
INLINE_HEADER int
i_used_to_be_a_crufty_macro_but_look_at_me_now (int arg)
{
    ...
}

```

The declaration right before the definition is to prevent warnings when compiling with `gcc -Wmissing-declarations`. I consider issuing this warning for inline functions a `gcc` bug, but the `gcc` maintainers disagree.

Every header which contains inline functions, either directly by using `INLINE_HEADER` or indirectly by using `DECLARE_LRECORD` must be added to `'inline.c'`'s includes to make the optimization described above work. (Optimization note: if all `INLINE_HEADER` functions are in fact inlined in all translation units, then the linker can just discard `inline.o`, since it contains only unreferenced code).

To get started debugging XEmacs, take a look at the `‘.gdbinit’` and `‘.dbxrc’` files in the `‘src’` directory. See the section in the XEmacs FAQ on How to Debug an XEmacs problem with a debugger.

After making source code changes, run `make check` to ensure that you haven't introduced any regressions. If you want to make xemacs more reliable, please improve the test suite in `‘tests/automated’`.

Did you make sure you didn't introduce any new compiler warnings?

Before submitting a patch, please try compiling at least once with

```
configure --with-mule --use-union-type --error-checking=all
```

Here are things to know when you create a new source file:

- All `‘.c’` files should `#include <config.h>` first. Almost all `‘.c’` files should `#include "lisp.h"` second.
- Generated header files should be included using the `#include <...>` syntax, not the `#include "..."` syntax. The generated headers are:

```
‘config.h sheap-adjust.h paths.h Emacs.ad.h’
```

The basic rule is that you should assume builds using `--srcdir` and the `#include <...>` syntax needs to be used when the to-be-included generated file is in a potentially different directory *at compile time*. The non-obvious C rule is that `#include "..."` means to search for the included file in the same directory as the including file, *not* in the current directory.

- Header files should *not* include `<config.h>` and `"lisp.h"`. It is the responsibility of the `‘.c’` files that use it to do so.

Here is a checklist of things to do when creating a new lisp object type named `foo`:

1. create `foo.h`
2. create `foo.c`
3. add definitions of `syms_of_foo`, etc. to `‘foo.c’`
4. add declarations of `syms_of_foo`, etc. to `‘symsinit.h’`
5. add calls to `syms_of_foo`, etc. to `‘emacs.c’`
6. add definitions of macros like `CHECK_FOO` and `FOO` to `‘foo.h’`
7. add the new type index to `enum lrecord_type`
8. add a `DEFINE_LRECORD_IMPLEMENTATION` call to `‘foo.c’`
9. add an `INIT_LRECORD_IMPLEMENTATION` call to `syms_of_foo.c`

## 11 Regression Testing XEmacs

The source directory ‘`tests/automated`’ contains XEmacs’ automated test suite. The usual way of running all the tests is running `make check` from the top-level build directory.

The test suite is unfinished and it’s still lacking some essential features. It is nevertheless recommended that you run the tests to confirm that XEmacs behaves correctly.

If you want to run a specific test case, you can do it from the command-line like this:

```
$ xemacs -batch -l test-harness.elc -f batch-test-emacs TEST-FILE
```

If a test fails and you need more information, you can run the test suite interactively by loading ‘`test-harness.el`’ into a running XEmacs and typing `M-x test-emacs-test-file RET<filename> RET`. You will see a log of passed and failed tests, which should allow you to investigate the source of the error and ultimately fix the bug. If you are not capable of, or don’t have time for, debugging it yourself, please do report the failures using `M-x report-emacs-bug` or `M-x build-report`.

**test-emacs-test-file** `le` [Command]

Runs the tests in `le`. ‘`test-harness.el`’ must be loaded. Defines all the macros described in this node, and undefines them when done.

Adding a new test file is trivial: just create a new file here and it will be run. There is no need to byte-compile any of the files in this directory—the test-harness will take care of any necessary byte-compilation.

Look at the existing test cases for the examples of coding test cases. It all boils down to your imagination and judicious use of the macros `Assert`, `Check-Error`, `Check-Error-Message`, and `Check-Message`. Note that all of these macros are defined only for the duration of the test: they do not exist in the global environment.

**Assert** `expr` [Macro]

Check that `expr` is non-nil at this point in the test.

**Check-Error** `expected-error body` [Macro]

Check that execution of `body` causes `expected-error` to be signaled. `body` is a `progn`-like body, and may contain several expressions. `expected-error` is a symbol defined as an error by `define-error`.

**Check-Error-Message** `expected-error expected-error-regexp body` [Macro]

Check that execution of `body` causes `expected-error` to be signaled, and generate a message matching `expected-error-regexp`. `body` is a `progn`-like body, and may contain several expressions. `expected-error` is a symbol defined as an error by `define-error`.

**Check-Message** `expected-message body` [Macro]

Check that execution of `body` causes `expected-message` to be generated (using `message` or a similar function). `body` is a `progn`-like body, and may contain several expressions.

Here’s a simple example checking case-sensitive and case-insensitive comparisons from ‘`case-tests.el`’.

```
(with-temp-buffer
  (insert "Test Buffer")
  (let ((case-fold-search t))
    (goto-char (point-min))
    (Assert (eq (search-forward "test buffer" nil t) 12))
    (goto-char (point-min))
    (Assert (eq (search-forward "Test buffer" nil t) 12))
    (goto-char (point-min))
    (Assert (eq (search-forward "Test Buffer" nil t) 12)))

  (setq case-fold-search nil)
  (goto-char (point-min))
  (Assert (not (search-forward "test buffer" nil t)))
  (goto-char (point-min))
  (Assert (not (search-forward "Test buffer" nil t)))
  (goto-char (point-min))
  (Assert (eq (search-forward "Test Buffer" nil t) 12))))
```

This example could be saved in a file in ‘tests/automated’, and it would constitute a complete test, automatically executed when you run **make check** after building XEmacs. More complex tests may require substantial temporary scaffolding to create the environment that elicits the bugs, but the top-level ‘Makefile’ and ‘test-harness.el’ handle the running and collection of results from the **Assert**, **Check-Error**, **Check-Error-Message**, and **Check-Message** macros.

Don’t suppress tests just because they’re due to known bugs not yet fixed—use the **Known-Bug-Expect-Failure** wrapper macro to mark them.

**Known-Bug-Expect-Failure** *body* [Macro]

Arrange for failing tests in **body** to generate messages prefixed with "KNOWN BUG:" instead of "FAIL:". **body** is a **progn**-like body, and may contain several tests.

A lot of the tests we run push limits; suppress Ebola warning messages with the **Ignore-Ebola** wrapper macro.

**Ignore-Ebola** *body* [Macro]

Suppress Ebola warning messages while running tests in **body**. **body** is a **progn**-like body, and may contain several tests.

Both macros are defined temporarily within the test function. Simple examples:

```
;; Apparently Ignore-Ebola is a solution with no problem to address.
;; There are no examples in 21.5, anyway.
```

```
;; from regexp-tests.el
(Known-Bug-Expect-Failure
  (Assert (not (string-match "\\b" "")))
  (Assert (not (string-match " \\b" " "))))
```

In general, you should avoid using functionality from packages in your tests, because you can’t be sure that everyone will have the required package. However, if you’ve got a test

that works, by all means add it. Simply wrap the test in an appropriate test, add a notice that the test was skipped, and update the `skipped-test-reasons` hashtable. The wrapper macro `Skip-Test-Unless` is provided to handle common cases.

**skipped-test-reasons** [Variable]

Hash table counting the number of times a particular reason is given for skipping tests. This is only defined within `test-emacs-test-file`.

**Skip-Test-Unless** `prerequisite reason description body` [Macro]

`prerequisite` is usually a feature test (`featurep`, `boundp`, `fboundp`). `reason` is a string describing the prerequisite; it must be unique because it is used as a hash key in a table of reasons for skipping tests. `description` describes the tests being skipped, for the test result summary. `body` is a `progn`-like body, and may contain several tests.

`Skip-Test-Unless` is defined temporarily within the test function. Here's an example of usage from `'syntax-tests.el'`:

```
;; Test forward-comment at buffer boundaries
(with-temp-buffer
  ;; try to use exactly what you need: featurep, boundp, fboundp
  (Skip-Test-Unless (fboundp 'c-mode)
                    "c-mode unavailable"
                    "comment and parse-partial-sexp tests"
    ;; and here's the test code
    (c-mode)
    (insert "// comment\n")
    (forward-comment -2)
    (Assert (eq (point) (point-min)))
    (let ((point (point)))
      (insert "/* comment */")
      (goto-char point)
      (forward-comment 2)
      (Assert (eq (point) (point-max)))
      (parse-partial-sexp point (point-max))))))
```

`Skip-Test-Unless` is intended for use with features that are normally present in typical configurations. For truly optional features, or tests that apply to one of several alternative implementations (eg, to GTK widgets, but not Athena, Motif, MS Windows, or Carbon), simply silently suppress the test if the feature is not available.

Here are a few general hints for writing tests.

1. Include related successful cases. Fixes often break something.
2. Use the `Known-Bug-Expect-Failure` macro to mark the cases you know are going to fail. We want to be able to distinguish between regressions and other unexpected failures, and cases that have been (partially) analyzed but not yet repaired.
3. Mark the bug with the date of report. An "Unfixed since yyyy-mm-dd" gloss for `Known-Bug-Expect-Failure` is planned to further increase developer embarrassment (== incentive to fix the bug), but until then at least put a comment about the date so we can easily see when it was first reported.

4. It's a matter of your judgement, but you should often use generic tests (*e.g.*, `eq`) instead of more specific tests (`=` for numbers) even though you know that arguments "should" be of correct type. That is, if the functions used can return generic objects (typically `nil`), as well as some more specific type that will be returned on success. We don't want failures of those assertions reported as "other failures" (a `wrong-type-arg` signal, rather than a null return), we want them reported as "assertion failures."

One example is a test that tests `(= (string-match this that) 0)`, expecting a successful match. Now suppose `string-match` is broken such that the match fails. Then it will return `nil`, and `=` will signal "wrong-type-argument, number-char-or-marker-p, nil", generating an "other failure" in the report. But this should be reported as an assertion failure (the test failed in a foreseeable way), rather than something else (we don't know what happened because XEmacs is broken in a way that we weren't trying to test!)

## 12 CVS Techniques

### 12.1 Merging a Branch into the Trunk

1. If you haven't already done a merge, you will be merging from the branch point; otherwise you'll be merging from the last merge point, which should be marked by a tag, e.g. 'last-sync-ben-mule-21-5'. In the former case, create the last-sync tag, e.g.

```
crw rtag -r ben-mule-21-5-bp last-sync-ben-mule-21-5 xemacs
```

(You did create a branch point tag when you created the branch, didn't you?)

2. Check everything in on your branch.
3. Tag your branch with a pre-sync tag, e.g.

```
crw rtag -r ben-mule-21-5 ben-mule-21-5-pre-feb-20-2002-sync xemacs
```

Note, you need to use rtag and specify a version with '-r' (use '-r HEAD' if necessary) so that removed files are handled correctly in some obscure cases. See section 4.8 of the CVS manual.

4. Tag the trunk so you have a stable place to merge up to in case people are asynchronously committing to the trunk, e.g.

```
crw rtag -r HEAD main-branch-ben-mule-21-5-syncpoint-feb-20-2002 xemacs
```

```
crw rtag -F -r main-branch-ben-mule-21-5-syncpoint-feb-20-2002 next-sync-ben-mule-
```

Use -F in the second case because the name might already exist, e.g. if you've already done a merge. We make two tags because one is a permanent mark indicating a syncpoint when merging, and the other is a symbolic tag to make other operations easier.

5. Make a backup of your source tree (not totally necessary but useful for reference and peace of mind): Move one level up from the top directory of your branch and do, e.g.

```
cp -a mule mule-backup-2-23-02
```

6. Now, we're ready to merge! Make sure you're in the top directory of your branch and do, e.g.

```
cvs update -j last-sync-ben-mule-21-5 -j next-sync-ben-mule-21-5
```

7. Fix all merge conflicts. Get the sucker to compile and run.
8. Tag your branch with a post-sync tag, e.g.

```
crw rtag -r ben-mule-21-5 ben-mule-21-5-post-feb-20-2002-sync xemacs
```

9. Update the last-sync tag, e.g.

```
crw rtag -F -r next-sync-ben-mule-21-5 last-sync-ben-mule-21-5 xemacs
```



## 13 A Summary of the Various XEmacs Modules

This is accurate as of XEmacs 20.0.

### 13.1 Low-Level Modules

`config.h`

This is automatically generated from ‘`config.h.in`’ based on the results of configure tests and user-selected optional features and contains preprocessor definitions specifying the nature of the environment in which XEmacs is being compiled.

`paths.h`

This is automatically generated from ‘`paths.h.in`’ based on supplied configure values, and allows for non-standard installed configurations of the XEmacs directories. It’s currently broken, though.

`emacs.c`

`signal.c`

‘`emacs.c`’ contains `main()` and other code that performs the most basic environment initializations and handles shutting down the XEmacs process (this includes `kill-emacs`, the normal way that XEmacs is exited; `dump-emacs`, which is used during the build process to write out the XEmacs executable; `run-emacs-from-temacs`, which can be used to start XEmacs directly when `temacs` has finished loading all the Lisp code; and emergency code to handle crashes [XEmacs tries to auto-save all files before it crashes]).

Low-level code that directly interacts with the Unix signal mechanism, however, is in ‘`signal.c`’. Note that this code does not handle system dependencies in interfacing to signals; that is handled using the ‘`sysignal.h`’ header file, described in section J below.

`unexaix.c`

`unexalpha.c`

`unexapollo.c`

`unexconvex.c`

`unexec.c`

`unexelf.c`

`unexelfsgi.c`

`unexencap.c`

`unexenix.c`

`unexfreebsd.c`

`unefx2800.c`

`unexhp9k3.c`

`unexhp9k800.c`

`unexmips.c`

`unexnext.c`

`unexsol2.c`

`unexsunos4.c`

These modules contain code dumping out the XEmacs executable on various different systems. (This process is highly machine-specific and requires intimate knowledge of the executable format and the memory map of the process.) Only one of these modules is actually used; this is chosen by ‘`configure`’.

```
ecrt0.c
lastfile.c
pre-crt0.c
```

These modules are used in conjunction with the dump mechanism. On some systems, an alternative version of the C startup code (the actual code that receives control from the operating system when the process is started, and which calls `main()`) is required so that the dumping process works properly; `crt0.c` provides this.

`pre-crt0.c` and `lastfile.c` should be the very first and very last file linked, respectively. (Actually, this is not really true. `lastfile.c` should be after all Emacs modules whose initialized data should be made constant, and before all other Emacs files and all libraries. In particular, the allocation modules `gmalloc.c`, `alloca.c`, etc. are normally placed past `lastfile.c`, and all of the files that implement Xt widget classes *must* be placed after `lastfile.c` because they contain various structures that must be statically initialized and into which Xt writes at various times.) `pre-crt0.c` and `lastfile.c` contain exported symbols that are used to determine the start and end of XEmacs' initialized data space when dumping.

```
alloca.c
free-hook.c
getpagesize.h
gmalloc.c
malloc.c
mem-limits.h
ralloc.c
vm-limit.c
```

These handle basic C allocation of memory. `alloca.c` is an emulation of the stack allocation function `alloca()` on machines that lack this. (XEmacs makes extensive use of `alloca()` in its code.)

`gmalloc.c` and `malloc.c` are two implementations of the standard C functions `malloc()`, `realloc()` and `free()`. They are often used in place of the standard system-provided `malloc()` because they usually provide a much faster implementation, at the expense of additional memory use. `gmalloc.c` is a newer implementation that is much more memory-efficient for large allocations than `malloc.c`, and should always be preferred if it works. (At one point, `gmalloc.c` didn't work on some systems where `malloc.c` worked; but this should be fixed now.)

`ralloc.c` is the **relocating allocator**. It provides functions similar to `malloc()`, `realloc()` and `free()` that allocate memory that can be dynamically relocated in memory. The advantage of this is that allocated memory can be shuffled around to place all the free memory at the end of the heap, and the heap can then be shrunk, releasing the memory back to the operating system. The use of this can be controlled with the configure option `--rel-alloc`; if enabled, memory allocated for buffers will be relocatable, so that if a very large file is visited and the buffer is later killed, the memory can be released to the operating system. (The disadvantage of this mechanism is that it can be very slow. On systems with the `mmap()` system call, the XEmacs version of `ralloc.c` uses this to move memory around without actually having to block-copy it, which can speed things up; but it can still cause noticeable performance degradation.)

On Linux systems using ‘glibc 2’, these strategies are built in to the so-called “Doug Lea malloc.” See, for example, Doug Lea’s home page, especially “[A Memory Allocator](#)”. The source file, ‘malloc.c’ (available at the same place) is copiously (and usefully!) commented. [Wolfram Gloger’s home page](#) may also be useful.

‘free-hook.c’ contains some debugging functions for checking for invalid arguments to `free()`.

‘vm-limit.c’ contains some functions that warn the user when memory is getting low. These are callback functions that are called by ‘gmalloc.c’ and ‘malloc.c’ at appropriate times.

‘getpagesize.h’ provides a uniform interface for retrieving the size of a page in virtual memory. ‘mem-limits.h’ provides a uniform interface for retrieving the total amount of available virtual memory. Both are similar in spirit to the ‘sys\*.h’ files described in section J, below.

```
blocktype.c
blocktype.h
dynarr.c
```

These implement a couple of basic C data types to facilitate memory allocation. The `Blocktype` type efficiently manages the allocation of fixed-size blocks by minimizing the number of times that `malloc()` and `free()` are called. It allocates memory in large chunks, subdivides the chunks into blocks of the proper size, and returns the blocks as requested. When blocks are freed, they are placed onto a linked list, so they can be efficiently reused. This data type is not much used in XEmacs currently, because it’s a fairly new addition.

The `Dynarr` type implements a **dynamic array**, which is similar to a standard C array but has no fixed limit on the number of elements it can contain. Dynamic arrays can hold elements of any type, and when you add a new element, the array automatically resizes itself if it isn’t big enough. Dynarrs are extensively used in the redisplay mechanism.

```
inline.c
```

This module is used in connection with inline functions (available in some compilers). Often, inline functions need to have a corresponding non-inline function that does the same thing. This module is where they reside. It contains no actual code, but defines some special flags that cause inline functions defined in header files to be rendered as actual functions. It then includes all header files that contain any inline function definitions, so that each one gets a real function equivalent.

```
debug.c
debug.h
```

These functions provide a system for doing internal consistency checks during code development. This system is not currently used; instead the simpler `assert()` macro is used along with the various checks provided by the ‘--error-check-\*’ configuration options.

```
universe.h
```

This is not currently used.

## 13.2 Basic Lisp Modules

```
lisp-disunion.h
lisp-union.h
```

```

lisp.h
lrecord.h
symsinit.h

```

These are the basic header files for all XEmacs modules. Each module includes ‘`lisp.h`’, which brings the other header files in. ‘`lisp.h`’ contains the definitions of the structures and extractor and constructor macros for the basic Lisp objects and various other basic definitions for the Lisp environment, as well as some general-purpose definitions (e.g. `min()` and `max()`). ‘`lisp.h`’ includes either ‘`lisp-disunion.h`’ or ‘`lisp-union.h`’, depending on whether `USE_UNION_TYPE` is defined. These files define the typedef of the Lisp object itself (as described above) and the low-level macros that hide the actual implementation of the Lisp object. All extractor and constructor macros for particular types of Lisp objects are defined in terms of these low-level macros.

As a general rule, all typedefs should go into the typedefs section of ‘`lisp.h`’ rather than into a module-specific header file even if the structure is defined elsewhere. This allows function prototypes that use the typedef to be placed into other header files. Forward structure declarations (i.e. a simple declaration like `struct foo;` where the structure itself is defined elsewhere) should be placed into the typedefs section as necessary.

‘`lrecord.h`’ contains the basic structures and macros that implement all record-type Lisp objects—i.e. all objects whose type is a field in their C structure, which includes all objects except the few most basic ones.

‘`lisp.h`’ contains prototypes for most of the exported functions in the various modules. Lisp primitives defined using `DEFUN` that need to be called by C code should be declared using `EXFUN`. Other function prototypes should be placed either into the appropriate section of `lisp.h`, or into a module-specific header file, depending on how general-purpose the function is and whether it has special-purpose argument types requiring definitions not in ‘`lisp.h`’.) All initialization functions are prototyped in ‘`symsinit.h`’.

```

alloc.c

```

The large module ‘`alloc.c`’ implements all of the basic allocation and garbage collection for Lisp objects. The most commonly used Lisp objects are allocated in chunks, similar to the Blocktype data type described above; others are allocated in individually `malloc()`ed blocks. This module provides the foundation on which all other aspects of the Lisp environment sit, and is the first module initialized at startup.

Note that ‘`alloc.c`’ provides a series of generic functions that are not dependent on any particular object type, and interfaces to particular types of objects using a standardized interface of type-specific methods. This scheme is a fundamental principle of object-oriented programming and is heavily used throughout XEmacs. The great advantage of this is that it allows for a clean separation of functionality into different modules—new classes of Lisp objects, new event interfaces, new device types, new stream interfaces, etc. can be added transparently without affecting code anywhere else in XEmacs. Because the different subsystems are divided into general and specific code, adding a new subtype within a subsystem will in general not require changes to the generic subsystem code or affect any of the other subtypes in the subsystem; this provides a great deal of robustness to the XEmacs code.

```

eval.c
backtrace.h

```

This module contains all of the functions to handle the flow of control. This includes the mechanisms of defining functions, calling functions, traversing stack frames, and binding variables; the control primitives and other special forms such as `while`, `if`, `eval`, `let`, `and`, `or`, `progn`, etc.; handling of non-local exits, `unwind-protects`, and exception handlers; entering the debugger; methods for the `subr` Lisp object type; etc. It does *not* include the `read` function, the `print` function, or the handling of symbols and obarrays.

‘`backtrace.h`’ contains some structures related to stack frames and the flow of control.

`lread.c`

This module implements the Lisp reader and the `read` function, which converts text into Lisp objects, according to the read syntax of the objects, as described above. This is similar to the parser that is a part of all compilers.

`print.c`

This module implements the Lisp print mechanism and the `print` function and related functions. This is the inverse of the Lisp reader – it converts Lisp objects to a printed, textual representation. (Hopefully something that can be read back in using `read` to get an equivalent object.)

`general.c`

`symbols.c`

`symeval.h`

‘`symbols.c`’ implements the handling of symbols, obarrays, and retrieving the values of symbols. Much of the code is devoted to handling the special **symbol-value-magic** objects that define special types of variables—this includes `buffer-local` variables, variable aliases, variables that forward into C variables, etc. This module is initialized extremely early (right after ‘`alloc.c`’), because it is here that the basic symbols `t` and `nil` are created, and those symbols are used everywhere throughout XEmacs.

‘`symeval.h`’ contains the definitions of symbol structures and the `DEFVAR_LISP()` and related macros for declaring variables.

`data.c`

`floatfns.c`

`fns.c`

These modules implement the methods and standard Lisp primitives for all the basic Lisp object types other than symbols (which are described above). ‘`data.c`’ contains all the predicates (primitives that return whether an object is of a particular type); the integer arithmetic functions; and the basic accessor and mutator primitives for the various object types. ‘`fns.c`’ contains all the standard predicates for working with sequences (where, abstractly speaking, a sequence is an ordered set of objects, and can be represented by a list, string, vector, or bit-vector); it also contains `equal`, perhaps on the grounds that bulk of the operation of `equal` is comparing sequences. ‘`floatfns.c`’ contains methods and primitives for floats and floating-point arithmetic.

`bytecode.c`

`bytecode.h`

‘`bytecode.c`’ implements the byte-code interpreter and compiled-function objects, and ‘`bytecode.h`’ contains associated structures. Note that the byte-code *compiler* is written in Lisp.

### 13.3 Modules for Standard Editing Operations

```
buffer.c
buffer.h
bufslots.h
```

‘`buffer.c`’ implements the `buffer` Lisp object type. This includes functions that create and destroy buffers; retrieve buffers by name or by other properties; manipulate lists of buffers (remember that buffers are permanent objects and stored in various ordered lists); retrieve or change buffer properties; etc. It also contains the definitions of all the built-in buffer-local variables (which can be viewed as buffer properties). It does *not* contain code to manipulate buffer-local variables (that’s in ‘`symbols.c`’, described above); or code to manipulate the text in a buffer.

‘`buffer.h`’ defines the structures associated with a buffer and the various macros for retrieving text from a buffer and special buffer positions (e.g. `point`, the default location for text insertion). It also contains macros for working with buffer positions and converting between their representations as character offsets and as byte offsets (under MULE, they are different, because characters can be multi-byte). It is one of the largest header files.

‘`bufslots.h`’ defines the fields in the buffer structure that correspond to the built-in buffer-local variables. It is its own header file because it is included many times in ‘`buffer.c`’, as a way of iterating over all the built-in buffer-local variables.

```
insdel.c
insdel.h
```

‘`insdel.c`’ contains low-level functions for inserting and deleting text in a buffer, keeping track of changed regions for use by redisplay, and calling any before-change and after-change functions that may have been registered for the buffer. It also contains the actual functions that convert between byte offsets and character offsets.

‘`insdel.h`’ contains associated headers.

```
marker.c
```

This module implements the `marker` Lisp object type, which conceptually is a pointer to a text position in a buffer that moves around as text is inserted and deleted, so as to remain in the same relative position. This module doesn’t actually move the markers around – that’s handled in ‘`insdel.c`’. This module just creates them and implements the primitives for working with them. As markers are simple objects, this does not entail much.

Note that the standard arithmetic primitives (e.g. `+`) accept markers in place of integers and automatically substitute the value of `marker-position` for the marker, i.e. an integer describing the current buffer position of the marker.

```
extents.c
extents.h
```

This module implements the `extent` Lisp object type, which is like a marker that works over a range of text rather than a single position. Extents are also much more complex and powerful than markers and have a more efficient (and more algorithmically complex) implementation. The implementation is described in detail in comments in ‘`extents.c`’.

The code in ‘`extents.c`’ works closely with ‘`insdel.c`’ so that extents are properly moved around as text is inserted and deleted. There is also code in ‘`extents.c`’ that provides information needed by the redisplay mechanism for efficient operation. (Remember that

extents can have display properties that affect [sometimes drastically, as in the `invisible` property] the display of the text they cover.)

`editfns.c`

`'editfns.c'` contains the standard Lisp primitives for working with a buffer's text, and calls the low-level functions in `'insdel.c'`. It also contains primitives for working with `point` (the default buffer insertion location).

`'editfns.c'` also contains functions for retrieving various characteristics from the external environment: the current time, the process ID of the running XEmacs process, the name of the user who ran this XEmacs process, etc. It's not clear why this code is in `'editfns.c'`.

`callint.c`

`cmds.c`

`commands.h`

These modules implement the basic **interactive** commands, i.e. user-callable functions. Commands, as opposed to other functions, have special ways of getting their parameters interactively (by querying the user), as opposed to having them passed in a normal function invocation. Many commands are not really meant to be called from other Lisp functions, because they modify global state in a way that's often undesired as part of other Lisp functions.

`'callint.c'` implements the mechanism for querying the user for parameters and calling interactive commands. The bulk of this module is code that parses the interactive spec that is supplied with an interactive command.

`'cmds.c'` implements the basic, most commonly used editing commands: commands to move around the current buffer and insert and delete characters. These commands are implemented using the Lisp primitives defined in `'editfns.c'`.

`'commands.h'` contains associated structure definitions and prototypes.

`regex.c`

`regex.h`

`search.c`

`'search.c'` implements the Lisp primitives for searching for text in a buffer, and some of the low-level algorithms for doing this. In particular, the fast fixed-string Boyer-Moore search algorithm is implemented in `'search.c'`. The low-level algorithms for doing regular-expression searching, however, are implemented in `'regex.c'` and `'regex.h'`. These two modules are largely independent of XEmacs, and are similar to (and based upon) the regular-expression routines used in `'grep'` and other GNU utilities.

`doprnt.c`

`'doprnt.c'` implements formatted-string processing, similar to `printf()` command in C.

`undo.c`

This module implements the undo mechanism for tracking buffer changes. Most of this could be implemented in Lisp.

## 13.4 Editor-Level Control Flow Modules

`event-Xt.c`

`event-msw.c`

```

event-stream.c
event-tty.c
events-mod.h
gpmevent.c
gpmevent.h
events.c
events.h

```

These implement the handling of events (user input and other system notifications).

‘`events.c`’ and ‘`events.h`’ define the **event** Lisp object type and primitives for manipulating it.

‘`event-stream.c`’ implements the basic functions for working with event queues, dispatching an event by looking it up in relevant keymaps and such, and handling timeouts; this includes the primitives **next-event** and **dispatch-event**, as well as related primitives such as **sit-for**, **sleep-for**, and **accept-process-output**. (‘`event-stream.c`’ is one of the hairiest and trickiest modules in XEmacs. Beware! You can easily mess things up here.)

‘`event-Xt.c`’ and ‘`event-tty.c`’ implement the low-level interfaces onto retrieving events from Xt (the X toolkit) and from TTY’s (using `read()` and `select()`), respectively. The event interface enforces a clean separation between the specific code for interfacing with the operating system and the generic code for working with events, by defining an API of basic, low-level event methods; ‘`event-Xt.c`’ and ‘`event-tty.c`’ are two different implementations of this API. To add support for a new operating system (e.g. NeXTstep), one merely needs to provide another implementation of those API functions.

Note that the choice of whether to use ‘`event-Xt.c`’ or ‘`event-tty.c`’ is made at compile time! Or at the very latest, it is made at startup time. ‘`event-Xt.c`’ handles events for *both* X and TTY frames; ‘`event-tty.c`’ is only used when X support is not compiled into XEmacs. The reason for this is that there is only one event loop in XEmacs: thus, it needs to be able to receive events from all different kinds of frames.

```

keymap.c
keymap.h

```

‘`keymap.c`’ and ‘`keymap.h`’ define the **keymap** Lisp object type and associated methods and primitives. (Remember that keymaps are objects that associate event descriptions with functions to be called to “execute” those events; **dispatch-event** looks up events in the relevant keymaps.)

```

cmdloop.c

```

‘`cmdloop.c`’ contains functions that implement the actual editor command loop—i.e. the event loop that cyclically retrieves and dispatches events. This code is also rather tricky, just like ‘`event-stream.c`’.

```

macros.c
macros.h

```

These two modules contain the basic code for defining keyboard macros. These functions don’t actually do much; most of the code that handles keyboard macros is mixed in with the event-handling code in ‘`event-stream.c`’.

```

minibuf.c

```

This contains some miscellaneous code related to the minibuffer (most of the minibuffer code was moved into Lisp by Richard Mlynarik). This includes the primitives for completion (although filename completion is in `'dired.c'`), the lowest-level interface to the minibuffer (if the command loop were cleaned up, this too could be in Lisp), and code for dealing with the echo area (this, too, was mostly moved into Lisp, and the only code remaining is code to call out to Lisp or provide simple bootstrapping implementations early in `temacs`, before the echo-area Lisp code is loaded).

## 13.5 Modules for the Basic Displayable Lisp Objects

```
console-msw.c
console-msw.h
console-stream.c
console-stream.h
console-tty.c
console-tty.h
console-x.c
console-x.h
console.c
console.h
```

These modules implement the `console` Lisp object type. A console contains multiple display devices, but only one keyboard and mouse. Most of the time, a console will contain exactly one device.

Consoles are the top of a lisp object inclusion hierarchy. Consoles contain devices, which contain frames, which contain windows.

```
device-msw.c
device-tty.c
device-x.c
device.c
device.h
```

These modules implement the `device` Lisp object type. This abstracts a particular screen or connection on which frames are displayed. As with Lisp objects, event interfaces, and other subsystems, the device code is separated into a generic component that contains a standardized interface (in the form of a set of methods) onto particular device types.

The device subsystem defines all the methods and provides method services for not only device operations but also for the frame, window, menubar, scrollbar, toolbar, and other displayable-object subsystems. The reason for this is that all of these subsystems have the same subtypes (X, TTY, NeXTstep, Microsoft Windows, etc.) as devices do.

```
frame-msw.c
frame-tty.c
frame-x.c
frame.c
frame.h
```

Each device contains one or more frames in which objects (e.g. text) are displayed. A frame corresponds to a window in the window system; usually this is a top-level window but it could potentially be one of a number of overlapping child windows within a top-level

window, using the MDI (Multiple Document Interface) protocol in Microsoft Windows or a similar scheme.

The ‘`frame-*`’ files implement the **frame** Lisp object type and provide the generic and device-type-specific operations on frames (e.g. raising, lowering, resizing, moving, etc.).

```

window.c
window.h

```

Each frame consists of one or more non-overlapping **windows** (better known as **panes** in standard window-system terminology) in which a buffer’s text can be displayed. Windows can also have scrollbars displayed around their edges.

‘`window.c`’ and ‘`window.h`’ implement the **window** Lisp object type and provide code to manage windows. Since windows have no associated resources in the window system (the window system knows only about the frame; no child windows or anything are used for XEmacs windows), there is no device-type-specific code here; all of that code is part of the redisplay mechanism or the code for particular object types such as scrollbars.

## 13.6 Modules for other Display-Related Lisp Objects

```

faces.c
faces.h

bitmaps.h
glyphs-eimage.c
glyphs-msw.c
glyphs-msw.h
glyphs-widget.c
glyphs-x.c
glyphs-x.h
glyphs.c
glyphs.h

objects-msw.c
objects-msw.h
objects-tty.c
objects-tty.h
objects-x.c
objects-x.h
objects.c
objects.h

menubar-msw.c
menubar-msw.h
menubar-x.c
menubar.c
menubar.h

scrollbar-msw.c
scrollbar-msw.h
scrollbar-x.c
scrollbar-x.h
scrollbar.c

```

```
scrollbar.h
toolbar-msw.c
toolbar-x.c
toolbar.c
toolbar.h
font-lock.c
```

This file provides C support for syntax highlighting—i.e. highlighting different syntactic constructs of a source file in different colors, for easy reading. The C support is provided so that this is fast.

```
dgif_lib.c
gif_err.c
gif_lib.h
gifalloc.c
```

These modules decode GIF-format image files, for use with glyphs. These files were removed due to Unisys patent infringement concerns.

## 13.7 Modules for the Redisplay Mechanism

```
redisplay-output.c
redisplay-msw.c
redisplay-tty.c
redisplay-x.c
redisplay.c
redisplay.h
```

These files provide the redisplay mechanism. As with many other subsystems in XEmacs, there is a clean separation between the general and device-specific support.

‘`redisplay.c`’ contains the bulk of the redisplay engine. These functions update the redisplay structures (which describe how the screen is to appear) to reflect any changes made to the state of any displayable objects (buffer, frame, window, etc.) since the last time that redisplay was called. These functions are highly optimized to avoid doing more work than necessary (since redisplay is called extremely often and is potentially a huge time sink), and depend heavily on notifications from the objects themselves that changes have occurred, so that redisplay doesn’t explicitly have to check each possible object. The redisplay mechanism also contains a great deal of caching to further speed things up; some of this caching is contained within the various displayable objects.

‘`redisplay-output.c`’ goes through the redisplay structures and converts them into calls to device-specific methods to actually output the screen changes.

‘`redisplay-x.c`’ and ‘`redisplay-tty.c`’ are two implementations of these redisplay output methods, for X frames and TTY frames, respectively.

```
indent.c
```

This module contains various functions and Lisp primitives for converting between buffer positions and screen positions. These functions call the redisplay mechanism to do most of the work, and then examine the redisplay structures to get the necessary information. This module needs work.

```
termcap.c
terminfo.c
tparam.c
```

These files contain functions for working with the termcap (BSD-style) and terminfo (System V style) databases of terminal capabilities and escape sequences, used when XEmacs is displaying in a TTY.

```
cm.c
cm.h
```

These files provide some miscellaneous TTY-output functions and should probably be merged into `'redisplay-tty.c'`.

## 13.8 Modules for Interfacing with the File System

```
lstream.c
lstream.h
```

These modules implement the **stream** Lisp object type. This is an internal-only Lisp object that implements a generic buffering stream. The idea is to provide a uniform interface onto all sources and sinks of data, including file descriptors, stdio streams, chunks of memory, Lisp buffers, Lisp strings, etc. That way, I/O functions can be written to the stream interface and can transparently handle all possible sources and sinks. (For example, the `read` function can read data from a file, a string, a buffer, or even a function that is called repeatedly to return data, without worrying about where the data is coming from or what-size chunks it is returned in.)

Note that in the C code, streams are called **lstreams** (for “Lisp streams”) to distinguish them from other kinds of streams, e.g. stdio streams and C++ I/O streams.

Similar to other subsystems in XEmacs, lstreams are separated into generic functions and a set of methods for the different types of lstreams. `'lstream.c'` provides implementations of many different types of streams; others are provided, e.g., in `'file-coding.c'`.

```
fileio.c
```

This implements the basic primitives for interfacing with the file system. This includes primitives for reading files into buffers, writing buffers into files, checking for the presence or accessibility of files, canonicalizing file names, etc. Note that these primitives are usually not invoked directly by the user: There is a great deal of higher-level Lisp code that implements the user commands such as `find-file` and `save-buffer`. This is similar to the distinction between the lower-level primitives in `'editfns.c'` and the higher-level user commands in `'commands.c'` and `'simple.el'`.

```
filelock.c
```

This file provides functions for detecting clashes between different processes (e.g. XEmacs and some external process, or two different XEmacs processes) modifying the same file. (XEmacs can optionally use the `'lock/'` subdirectory to provide a form of “locking” between different XEmacs processes.) This module is also used by the low-level functions in `'insdel.c'` to ensure that, if the first modification is being made to a buffer whose corresponding file has been externally modified, the user is made aware of this so that the buffer can be synched up with the external changes if necessary.

```
filemode.c
```

This file provides some miscellaneous functions that construct a ‘`rwxr-xr-x`’-type permissions string (as might appear in an ‘`ls`’-style directory listing) given the information returned by the `stat()` system call.

```

dired.c
ndir.h

```

These files implement the XEmacs interface to directory searching. This includes a number of primitives for determining the files in a directory and for doing filename completion. (Remember that generic completion is handled by a different mechanism, in ‘`minibuf.c`’.)

‘`ndir.h`’ is a header file used for the directory-searching emulation functions provided in ‘`sysdep.c`’ (see section J below), for systems that don’t provide any directory-searching functions. (On those systems, directories can be read directly as files, and parsed.)

```

realpath.c

```

This file provides an implementation of the `realpath()` function for expanding symbolic links, on systems that don’t implement it or have a broken implementation.

### 13.9 Modules for Other Aspects of the Lisp Interpreter and Object System

```

elhash.c
elhash.h
hash.c
hash.h

```

These files provide two implementations of hash tables. Files ‘`hash.c`’ and ‘`hash.h`’ provide a generic C implementation of hash tables which can stand independently of XEmacs. Files ‘`elhash.c`’ and ‘`elhash.h`’ provide a separate implementation of hash tables that can store only Lisp objects, and knows about Lispy things like garbage collection, and implement the **hash-table** Lisp object type.

```

specifier.c
specifier.h

```

This module implements the **specifier** Lisp object type. This is primarily used for displayable properties, and allows for values that are specific to a particular buffer, window, frame, device, or device class, as well as a default value existing. This is used, for example, to control the height of the horizontal scrollbar or the appearance of the **default**, **bold**, or other faces. The specifier object consists of a number of specifications, each of which maps from a buffer, window, etc. to a value. The function `specifier-instance` looks up a value given a window (from which a buffer, frame, and device can be derived).

```

chartab.c
chartab.h
casetab.c

```

‘`chartab.c`’ and ‘`chartab.h`’ implement the **char table** Lisp object type, which maps from characters or certain sorts of character ranges to Lisp objects. The implementation of this object type is optimized for the internal representation of characters. Char tables come in different types, which affect the allowed object types to which a character can be mapped and also dictate certain other properties of the char table.

‘`casetab.c`’ implements one sort of char table, the **case table** which maps characters to other characters of possibly different case. These are used by XEmacs to implement case-changing primitives and to do case-insensitive searching.

```

syntax.c
syntax.h

```

This module implements **syntax tables**, another sort of char table that maps characters into syntax classes that define the syntax of these characters (e.g. a parenthesis belongs to a class of ‘open’ characters that have corresponding ‘close’ characters and can be nested). This module also implements the Lisp **scanner**, a set of primitives for scanning over text based on syntax tables. This is used, for example, to find the matching parenthesis in a command such as `forward-sexp`, and by ‘`font-lock.c`’ to locate quoted strings, comments, etc.

Syntax codes are implemented as bitfields in an int. Bits 0-6 contain the syntax code itself, bit 7 is a special prefix flag used for Lisp, and bits 16-23 contain comment syntax flags. From the Lisp programmer’s point of view, there are 11 flags: 2 styles X 2 characters X {start, end} flags for two-character comment delimiters, 2 style flags for one-character comment delimiters, and the prefix flag.

Internally, however, the characters used in multi-character delimiters will have non-comment-character syntax classes (e.g., the ‘/’ in C’s ‘/\*’ comment-start delimiter has “punctuation” (here meaning “operator-like”) class in C modes). Thus in a mixed comment style, such as C++’s ‘//’ to end of line, is represented by giving ‘/’ the “punctuation” class and the “style b first character of start sequence” and “style b second character of start sequence” flags. The fact that class is *not* punctuation allows the syntax scanner to recognize that this is a multi-character delimiter. The ‘`newline`’ character is given (single-character) “comment-end” *class* and the “style b first character of end sequence” *flag*. The “comment-end” class allows the scanner to determine that no second character is needed to terminate the comment.

```

casefiddle.c

```

This module implements various Lisp primitives for upcasing, downcasing and capitalizing strings or regions of buffers.

```

rangetab.c

```

This module implements the **range table** Lisp object type, which provides for a mapping from ranges of integers to arbitrary Lisp objects.

```

opaque.c
opaque.h

```

This module implements the **opaque** Lisp object type, an internal-only Lisp object that encapsulates an arbitrary block of memory so that it can be managed by the Lisp allocation system. To create an opaque object, you call `make_opaque()`, passing a pointer to a block of memory. An object is created that is big enough to hold the memory, which is copied into the object’s storage. The object will then stick around as long as you keep pointers to it, after which it will be automatically reclaimed.

Opaque objects can also have an arbitrary **mark method** associated with them, in case the block of memory contains other Lisp objects that need to be marked for garbage-collection purposes. (If you need other object methods, such as a finalize method, you should just go ahead and create a new Lisp object type—it’s not hard.)

**abbrev.c**

This function provides a few primitives for doing dynamic abbreviation expansion. In XEmacs, most of the code for this has been moved into Lisp. Some C code remains for speed and because the primitive `self-insert-command` (which is executed for all self-inserting characters) hooks into the abbrev mechanism. (`self-insert-command` is itself in C only for speed.)

**doc.c**

This function provides primitives for retrieving the documentation strings of functions and variables. These documentation strings contain certain special markers that get dynamically expanded (e.g. a reverse-lookup is performed on some named functions to retrieve their current key bindings). Some documentation strings (in particular, for the built-in primitives and pre-loaded Lisp functions) are stored externally in a file ‘DOC’ in the ‘lib-src/’ directory and need to be fetched from that file. (Part of the build stage involves building this file, and another part involves constructing an index for this file and embedding it into the executable, so that the functions in ‘doc.c’ do not have to search the entire ‘DOC’ file to find the appropriate documentation string.)

**md5.c**

This function provides a Lisp primitive that implements the MD5 secure hashing scheme, used to create a large hash value of a string of data such that the data cannot be derived from the hash value. This is used for various security applications on the Internet.

## 13.10 Modules for Interfacing with the Operating System

**callproc.c****process.c****process.h**

These modules allow XEmacs to spawn and communicate with subprocesses and network connections.

‘callproc.c’ implements (through the `call-process` primitive) what are called **synchronous subprocesses**. This means that XEmacs runs a program, waits till it’s done, and retrieves its output. A typical example might be calling the ‘ls’ program to get a directory listing.

‘process.c’ and ‘process.h’ implement **asynchronous subprocesses**. This means that XEmacs starts a program and then continues normally, not waiting for the process to finish. Data can be sent to the process or retrieved from it as it’s running. This is used for the `shell` command (which provides a front end onto a shell program such as ‘csh’), the mail and news readers implemented in XEmacs, etc. The result of calling `start-process` to start a subprocess is a process object, a particular kind of object used to communicate with the subprocess. You can send data to the process by passing the process object and the data to `send-process`, and you can specify what happens to data retrieved from the process by setting properties of the process object. (When the process sends data, XEmacs receives a process event, which says that there is data ready. When `dispatch-event` is called on this event, it reads the data from the process and does something with it, as specified by the process object’s properties. Typically, this means inserting the data into a buffer or calling a function.) Another property of the process object is called the **sentinel**, which is a function that is called when the process terminates.

Process objects are also used for network connections (connections to a process running on another machine). Network connections are started with `open-network-stream` but otherwise work just like subprocesses.

```
sysdep.c
sysdep.h
```

These modules implement most of the low-level, messy operating-system interface code. This includes various device control (ioctl) operations for file descriptors, TTY's, pseudo-terminals, etc. (usually this stuff is fairly system-dependent; thus the name of this module), and emulation of standard library functions and system calls on systems that don't provide them or have broken versions.

```
sysdir.h
sysfile.h
sysfloat.h
sysproc.h
syspwd.h
syssignal.h
systime.h
systty.h
syswait.h
```

These header files provide consistent interfaces onto system-dependent header files and system calls. The idea is that, instead of including a standard header file like `<sys/param.h>` (which may or may not exist on various systems) or having to worry about whether all systems provide a particular preprocessor constant, or having to deal with the four different paradigms for manipulating signals, you just include the appropriate `'sys*.h'` header file, which includes all the right system header files, defines and missing preprocessor constants, provides a uniform interface onto system calls, etc.

`'sysdir.h'` provides a uniform interface onto directory-querying functions. (In some cases, this is in conjunction with emulation functions in `'sysdep.c'`.)

`'sysfile.h'` includes all the necessary header files for standard system calls (e.g. `read()`), ensures that all necessary `open()` and `stat()` preprocessor constants are defined, and possibly (usually) substitutes sugared versions of `read()`, `write()`, etc. that automatically restart interrupted I/O operations.

`'sysfloat.h'` includes the necessary header files for floating-point operations.

`'sysproc.h'` includes the necessary header files for calling `select()`, `fork()`, `execve()`, socket operations, and the like, and ensures that the `FD_*()` macros for descriptor-set manipulations are available.

`'syspwd.h'` includes the necessary header files for obtaining information from `/etc/passwd` (the functions are emulated under VMS).

`'syssignal.h'` includes the necessary header files for signal-handling and provides a uniform interface onto the different signal-handling and signal-blocking paradigms.

`'systime.h'` includes the necessary header files and provides uniform interfaces for retrieving the time of day, setting file access/modification times, getting the amount of time used by the XEmacs process, etc.

`'systty.h'` buffers against the infinitude of different ways of controlling TTY's.

'`syswait.h`' provides a uniform way of retrieving the exit status from a `wait()`ed-on process (some systems use a union, others use an int).

```
hpplay.c
libsst.c
libsst.h
libst.h
linuxplay.c
nas.c
sgisplay.c
sound.c
sunplay.c
```

These files implement the ability to play various sounds on some types of computers. You have to configure your XEmacs with sound support in order to get this capability.

'`sound.c`' provides the generic interface. It implements various Lisp primitives and variables that let you specify which sounds should be played in certain conditions. (The conditions are identified by symbols, which are passed to `ding` to make a sound. Various standard functions call this function at certain times; if sound support does not exist, a simple beep results.

'`sgisplay.c`', '`sunplay.c`', '`hpplay.c`', and '`linuxplay.c`' interface to the machine's speaker for various different kind of machines. This is called **native** sound.

'`nas.c`' interfaces to a computer somewhere else on the network using the NAS (Network Audio Server) protocol, playing sounds on that machine. This allows you to run XEmacs on a remote machine, with its display set to your local machine, and have the sounds be made on your local machine, provided that you have a NAS server running on your local machine.

'`libsst.c`', '`libsst.h`', and '`libst.h`' provide some additional functions for playing sound on a Sun SPARC but are not currently in use.

```
tooltalk.c
tooltalk.h
```

These two modules implement an interface to the ToolTalk protocol, which is an interprocess communication protocol implemented on some versions of Unix. ToolTalk is a high-level protocol that allows processes to register themselves as providers of particular services; other processes can then request a service without knowing or caring exactly who is providing the service. It is similar in spirit to the DDE protocol provided under Microsoft Windows. ToolTalk is a part of the new CDE (Common Desktop Environment) specification and is used to connect the parts of the SPARCWorks development environment.

```
getloadavg.c
```

This module provides the ability to retrieve the system's current load average. (The way to do this is highly system-specific, unfortunately, and requires a lot of special-case code.)

```
sunpro.c
```

This module provides a small amount of code used internally at Sun to keep statistics on the usage of XEmacs.

```
broken-sun.h
strcmp.c
```

```
strcpy.c
sunOS-fix.c
```

These files provide replacement functions and prototypes to fix numerous bugs in early releases of SunOS 4.1.

```
hftctl.c
```

This module provides some terminal-control code necessary on versions of AIX prior to 4.1.

## 13.11 Modules for Interfacing with X Windows

```
Emacs.ad.h
```

A file generated from ‘Emacs.ad’, which contains XEmacs-supplied fallback resources (so that XEmacs has pretty defaults).

```
EmacsFrame.c
EmacsFrame.h
EmacsFrameP.h
```

These modules implement an Xt widget class that encapsulates a frame. This is for ease in integrating with Xt. The EmacsFrame widget covers the entire X window except for the menubar; the scrollbars are positioned on top of the EmacsFrame widget.

**Warning:** Abandon hope, all ye who enter here. This code took an ungodly amount of time to get right, and is likely to fall apart mercilessly at the slightest change. Such is life under Xt.

```
EmacsManager.c
EmacsManager.h
EmacsManagerP.h
```

These modules implement a simple Xt manager (i.e. composite) widget class that simply lets its children set whatever geometry they want. It’s amazing that Xt doesn’t provide this standardly, but on second thought, it makes sense, considering how amazingly broken Xt is.

```
EmacsShell-sub.c
EmacsShell.c
EmacsShell.h
EmacsShellP.h
```

These modules implement two Xt widget classes that are subclasses of the TopLevelShell and TransientShell classes. This is necessary to deal with more brokenness that Xt has sadistically thrust onto the backs of developers.

```
xgccache.c
xgccache.h
```

These modules provide functions for maintenance and caching of GC’s (graphics contexts) under the X Window System. This code is junky and needs to be rewritten.

```
select-msw.c
select-x.c
select.c
select.h
```

This module provides an interface to the X Window System's concept of **selections**, the standard way for X applications to communicate with each other.

```
xintrinsic.h
xintrinsicp.h
xmmanagerp.h
xmprimitivep.h
```

These header files are similar in spirit to the 'sys\*.h' files and buffer against different implementations of Xt and Motif.

- 'xintrinsic.h' should be included in place of '<Intrinsic.h>'.
- 'xintrinsicp.h' should be included in place of '<IntrinsicP.h>'.
- 'xmmanagerp.h' should be included in place of '<XmManagerP.h>'.
- 'xmprimitivep.h' should be included in place of '<XmPrimitiveP.h>'.

```
xmu.c
xmu.h
```

These files provide an emulation of the Xmu library for those systems (i.e. HP-UX) that don't provide it as a standard part of X.

```
ExternalClient-Xlib.c
ExternalClient.c
ExternalClient.h
ExternalClientP.h
ExternalShell.c
ExternalShell.h
ExternalShellP.h
extw-Xlib.c
extw-Xlib.h
extw-Xt.c
extw-Xt.h
```

These files provide the **external widget** interface, which allows an XEmacs frame to appear as a widget in another application. To do this, you have to configure with '--external-widget'.

'ExternalShell\*' provides the server (XEmacs) side of the connection.

'ExternalClient\*' provides the client (other application) side of the connection. These files are not compiled into XEmacs but are compiled into libraries that are then linked into your application.

'extw-\*' is common code that is used for both the client and server.

Don't touch this code; something is liable to break if you do.

## 13.12 Modules for Internationalization

```
mule-canna.c
mule-ccl.c
mule-charset.c
mule-charset.h
file-coding.c
```

```
file-coding.h
mule-mcpath.c
mule-mcpath.h
mule-wnnfns.c
mule.c
```

These files implement the MULE (Asian-language) support. Note that MULE actually provides a general interface for all sorts of languages, not just Asian languages (although they are generally the most complicated to support). This code is still in beta.

‘`mule-charset.*`’ and ‘`file-coding.*`’ provide the heart of the XEmacs MULE support. ‘`mule-charset.*`’ implements the `charset` Lisp object type, which encapsulates a character set (an ordered one- or two-dimensional set of characters, such as US ASCII or JISX0208 Japanese Kanji).

‘`file-coding.*`’ implements the `coding-system` Lisp object type, which encapsulates a method of converting between different encodings. An encoding is a representation of a stream of characters, possibly from multiple character sets, using a stream of bytes or words, and defines (e.g.) which escape sequences are used to specify particular character sets, how the indices for a character are converted into bytes (sometimes this involves setting the high bit; sometimes complicated rearranging of the values takes place, as in the Shift-JIS encoding), etc.

‘`mule-ccl.c`’ provides the CCL (Code Conversion Language) interpreter. CCL is similar in spirit to Lisp byte code and is used to implement converters for custom encodings.

‘`mule-canna.c`’ and ‘`mule-wnnfns.c`’ implement interfaces to external programs used to implement the Canna and WNN input methods, respectively. This is currently in beta.

‘`mule-mcpath.c`’ provides some functions to allow for pathnames containing extended characters. This code is fragmentary, obsolete, and completely non-working. Instead, `pathname-coding-system` is used to specify conversions of names of files and directories. The standard C I/O functions like ‘`open()`’ are wrapped so that conversion occurs automatically.

‘`mule.c`’ provides a few miscellaneous things that should probably be elsewhere.

```
intl.c
```

This provides some miscellaneous internationalization code for implementing message translation and interfacing to the Ximp input method. None of this code is currently working.

```
iso-wide.h
```

This contains leftover code from an earlier implementation of Asian-language support, and is not currently used.

### 13.13 Modules for Regression Testing

```
test-harness.el
base64-tests.el
byte-compiler-tests.el
case-tests.el
ccl-tests.el
c-tests.el
```

```
database-tests.el
extent-tests.el
hash-table-tests.el
lisp-tests.el
md5-tests.el
mule-tests.el
regexp-tests.el
symbol-tests.el
syntax-tests.el
```

'test-harness.el' defines the macros `Assert`, `Check-Error`, `Check-Error-Message`, and `Check-Message`. The other files are test files, testing various XEmacs facilities.



## 14 Allocation of Objects in XEmacs Lisp

### 14.1 Introduction to Allocation

Emacs Lisp, like all Lisps, has garbage collection. This means that the programmer never has to explicitly free (destroy) an object; it happens automatically when the object becomes inaccessible. Most experts agree that garbage collection is a necessity in a modern, high-level language. Its omission from C stems from the fact that C was originally designed to be a nice abstract layer on top of assembly language, for writing kernels and basic system utilities rather than large applications.

Lisp objects can be created by any of a number of Lisp primitives. Most object types have one or a small number of basic primitives for creating objects. For conses, the basic primitive is `cons`; for vectors, the primitives are `make-vector` and `vector`; for symbols, the primitives are `make-symbol` and `intern`; etc. Some Lisp objects, especially those that are primarily used internally, have no corresponding Lisp primitives. Every Lisp object, though, has at least one C primitive for creating it.

Recall from section (VII) that a Lisp object, as stored in a 32-bit or 64-bit word, has a few tag bits, and a “value” that occupies the remainder of the bits. We can separate the different Lisp object types into three broad categories:

- (a) Those for whom the value directly represents the contents of the Lisp object. Only two types are in this category: integers and characters. No special allocation or garbage collection is necessary for such objects. Lisp objects of these types do not need to be GCPR0ed.

In the remaining two categories, the type is stored in the object itself. The tag for all such objects is the generic `lrecord` (`Lisp_Type_Record`) tag. The first bytes of the object’s structure are an integer (actually a `char`) characterising the object’s type and some flags, in particular the mark bit used for garbage collection. A structure describing the type is accessible thru the `lrecord_implementation_table` indexed with said integer. This structure includes the method pointers and a pointer to a string naming the type.

- (b) Those `lrecords` that are allocated in frob blocks (see above). This includes the objects that are most common and relatively small, and includes conses, strings, subrs, floats, compiled functions, symbols, extents, events, and markers. With the cleanup of frob blocks done in 19.12, it’s not terribly hard to add more objects to this category, but it’s a bit trickier than adding an object type to type (c) (esp. if the object needs a finalization method), and is not likely to save much space unless the object is small and there are many of them. (In fact, if there are very few of them, it might actually waste space.)
- (c) Those `lrecords` that are individually `malloc()`ed. These are called `lrecords`. All other types are in this category. Adding a new type to this category is comparatively easy, and all types added since 19.8 (when the current allocation scheme was devised, by Richard Mlynarik), with the exception of the character type, have been in this category.

Note that bit vectors are a bit of a special case. They are simple `lrecords` as in category (b), but are individually `malloc()`ed like vectors. You can basically view them as exactly

like vectors except that their type is stored in lrecord fashion rather than in directly-tagged fashion.

## 14.2 Garbage Collection

Garbage collection is simple in theory but tricky to implement. Emacs Lisp uses the oldest garbage collection method, called **mark and sweep**. Garbage collection begins by starting with all accessible locations (i.e. all variables and other slots where Lisp objects might occur) and recursively traversing all objects accessible from those slots, marking each one that is found. We then go through all of memory and free each object that is not marked, and unmarking each object that is marked. Note that “all of memory” means all currently allocated objects. Traversing all these objects means traversing all lrecord blocks, all vectors (which are chained in one big list), and all lrecords (which are likewise chained).

Garbage collection can be invoked explicitly by calling `garbage-collect` but is also called automatically by `eval`, once a certain amount of memory has been allocated since the last garbage collection (according to `gc-cons-threshold`).

## 14.3 GCPR0g

GCPR0ing is one of the ugliest and trickiest parts of Emacs internals. The basic idea is that whenever garbage collection occurs, all in-use objects must be reachable somehow or other from one of the roots of accessibility. The roots of accessibility are:

1. All objects that have been `staticpro()`d or `staticpro_nodump()`ed. This is used for any global C variables that hold Lisp objects. A call to `staticpro()` happens implicitly as a result of any symbols declared with `defsymbol()` and any variables declared with `DEFVAR_FOO()`. You need to explicitly call `staticpro()` (in the `vars_of_foo()` method of a module) for other global C variables holding Lisp objects. (This typically includes internal lists and such things.). Use `staticpro_nodump()` only in the rare cases when you do not want the pointed variable to be saved at dump time but rather recompute it at startup.

Note that `obarray` is one of the `staticpro()`d things. Therefore, all functions and variables get marked through this.

2. Any shadowed bindings that are sitting on the `specpdl` stack.
3. Any objects sitting in currently active (Lisp) stack frames, catches, and condition cases.
4. A couple of special-case places where active objects are located.
5. Anything currently marked with `GCPR0`.

Marking with `GCPR0` is necessary because some C functions (quite a lot, in fact), allocate objects during their operation. Quite frequently, there will be no other pointer to the object while the function is running, and if a garbage collection occurs and the object needs to be referenced again, bad things will happen. The solution is to mark those objects with `GCPR0`. Unfortunately this is easy to forget, and there is basically no way around this problem. Here are some rules, though:

1. For every `GCPR0n`, there have to be declarations of `struct gcpro gcpro1, gcpro2, etc.`
2. You *must* `UNGCPRO` anything that's `GCPR0`ed, and you *must not* `UNGCPRO` if you haven't `GCPR0`ed. Getting either of these wrong will lead to crashes, often in completely random places unrelated to where the problem lies.

3. The way this actually works is that all currently active `GCPROs` are chained through the `struct gcpro` local variables, with the variable `'gcprolist'` pointing to the head of the list and the `n`th local `gcpro` variable pointing to the first `gcpro` variable in the next enclosing stack frame. Each `GCPROed` thing is an lvalue, and the `struct gcpro` local variable contains a pointer to this lvalue. This is why things will mess up badly if you don't pair up the `GCPROs` and `UNGCPROs`—you will end up with `gcprolists` containing pointers to `struct gcpros` or local `Lisp_Object` variables in no-longer-active stack frames.
4. It is actually possible for a single `struct gcpro` to protect a contiguous array of any number of values, rather than just a single lvalue. To effect this, call `GCPRON` as usual on the first object in the array and then set `gcpron.nvars`.
5. **Strings are relocated.** What this means in practice is that the pointer obtained using `XSTRING_DATA()` is liable to change at any time, and you should never keep it around past any function call, or pass it as an argument to any function that might cause a garbage collection. This is why a number of functions accept either a “non-relocatable” `char *` pointer or a relocatable Lisp string, and only access the Lisp string's data at the very last minute. In some cases, you may end up having to `alloca()` some space and copy the string's data into it.
6. By convention, if you have to nest `GCPRO's`, use `NGCPRON` (along with `struct gcpro ngcpro1, ngcpro2, etc.`), `NNGCPRON`, etc. This avoids compiler warnings about shadowed locals.
7. It is *always* better to err on the side of extra `GCPROs` rather than too few. The extra cycles spent on this are almost never going to make a whit of difference in the speed of anything.
8. The general rule to follow is that caller, not callee, `GCPROs`. That is, you should not have to explicitly `GCPRO` any Lisp objects that are passed in as parameters.  
One exception from this rule is if you ever plan to change the parameter value, and store a new object in it. In that case, you *must* `GCPRO` the parameter, because otherwise the new object will not be protected.  
So, if you create any Lisp objects (remember, this happens in all sorts of circumstances, e.g. with `Fcons()`, etc.), you are responsible for `GCPROing` them, unless you are *absolutely sure* that there's no possibility that a garbage-collection can occur while you need to use the object. Even then, consider `GCPROing`.
9. A garbage collection can occur whenever anything calls `Feval`, or whenever a `QUIT` can occur where execution can continue past this. (Remember, this is almost anywhere.)
10. If you have the *least smidgeon of doubt* about whether you need to `GCPRO`, you should `GCPRO`.
11. Beware of `GCPROing` something that is uninitialized. If you have any shade of doubt about this, initialize all your variables to `Qnil`.
12. Be careful of traps, like calling `Fcons()` in the argument to another function. By the “caller protects” law, you should be `GCPROing` the newly-created cons, but you aren't. A certain number of functions that are commonly called on freshly created stuff (e.g. `nconc2()`, `Fsignal()`), break the “caller protects” law and go ahead and `GCPRO` their arguments so as to simplify things, but make sure and check if it's OK whenever doing something like this.

13. Once again, remember to **GCPR0**! Bugs resulting from insufficient **GCPR0**ing are intermittent and extremely difficult to track down, often showing up in crashes inside of `garbage-collect` or in weirdly corrupted objects or even in incorrect values in a totally different section of code.

If you don't understand whether to **GCPR0** in a particular instance, ask on the mailing lists. A general hint is that `prog1` is the canonical example.

Given the extremely error-prone nature of the **GCPR0** scheme, and the difficulties in tracking down, it should be considered a deficiency in the XEmacs code. A solution to this problem would involve implementing so-called **conservative** garbage collection for the C stack. That involves looking through all of stack memory and treating anything that looks like a reference to an object as a reference. This will result in a few objects not getting collected when they should, but it obviates the need for **GCPR0**ing, and allows garbage collection to happen at any point at all, such as during object allocation.

## 14.4 Garbage Collection - Step by Step

### 14.4.1 Invocation

The first thing that anyone should know about garbage collection is: when and how the garbage collector is invoked. One might think that this could happen every time new memory is allocated, e.g. new objects are created, but this is *not* the case. Instead, we have the following situation:

The entry point of any process of garbage collection is an invocation of the function `garbage_collect_1` in file `alloc.c`. The invocation can occur *explicitly* by calling the function `Fgarbage_collect` (in addition this function provides information about the freed memory), or can occur *implicitly* in four different situations:

1. In function `main_1` in file `emacs.c`. This function is called at each startup of xemacs. The garbage collection is invoked after all initial creations are completed, but only if a special internal error checking-constant `ERROR_CHECK_GC` is defined.
2. In function `disksave_object_finalization` in file `alloc.c`. The only purpose of this function is to clear the objects from memory which need not be stored with xemacs when we dump out an executable. This is only done by `Fdump_emacs` or by `Fdump_emacs_data` respectively (both in `emacs.c`). The actual clearing is accomplished by making these objects unreachable and starting a garbage collection. The function is only used while building xemacs.
3. In function `Feval` / `eval` in file `eval.c`. Each time the well known and often used function `eval` is called to evaluate a form, one of the first things that could happen, is a potential call of `garbage_collect_1`. There exist three global variables, `consing_since_gc` (counts the created cons-cells since the last garbage collection), `gc_cons_threshold` (a specified threshold after which a garbage collection occurs) and `always_gc`. If `always_gc` is set or if the threshold is exceeded, the garbage collection will start.
4. In function `Ffuncall` / `funcall` in file `eval.c`. This function evaluates calls of elisp functions and works according to `Feval`.

The upshot is that garbage collection can basically occur everywhere `Feval`, respectively `Ffuncall`, is used - either directly or through another function. Since calls to these two func-

tions are hidden in various other functions, many calls to `garbage_collect_1` are not obviously foreseeable, and therefore unexpected. Instances where they are used that are worth remembering are various elisp commands, as for example `or`, `and`, `if`, `cond`, `while`, `setq`, etc., miscellaneous `gui_item_...` functions, everything related to `eval` (`Feval_buffer`, `call0`, ...) and inside `Fsignal`. The latter is used to handle signals, as for example the ones raised by every `QUIT`-macro triggered after pressing `Ctrl-g`.

### 14.4.2 `garbage_collect_1`

We can now describe exactly what happens after the invocation takes place.

1. There are several cases in which the garbage collector is left immediately: when we are already garbage collecting (`gc_in_progress`), when the garbage collection is somehow forbidden (`gc_currently_forbidden`), when we are currently displaying something (`in_display`) or when we are preparing for the armageddon of the whole system (`preparing_for_armageddon`).
2. Next the correct frame in which to put all the output occurring during garbage collecting is determined. In order to be able to restore the old display's state after displaying the message, some data about the current cursor position has to be saved. The variables `pre_gc_cursor` and `cursor_changed` take care of that.
3. The state of `gc_currently_forbidden` must be restored after the garbage collection, no matter what happens during the process. We accomplish this by `record_unwind_protecting` the suitable function `restore_gc_inhibit` together with the current value of `gc_currently_forbidden`.
4. If we are concurrently running an interactive xemacs session, the next step is simply to show the garbage collector's cursor/message.
5. The following steps are the intrinsic steps of the garbage collector, therefore `gc_in_progress` is set.
6. For debugging purposes, it is possible to copy the current C stack frame. However, this seems to be a currently unused feature.
7. Before actually starting to go over all live objects, references to objects that are no longer used are pruned. We only have to do this for events (`clear_event_resource`) and for specifiers (`cleanup_specifiers`).
8. Now the mark phase begins and marks all accessible elements. In order to start from all slots that serve as roots of accessibility, the function `mark_object` is called for each root individually to go out from there to mark all reachable objects. All roots that are traversed are shown in their processed order:
  - all constant symbols and static variables that are registered via `staticpro` in the dynarr `staticpros`. See [Section 10.5 \[Adding Global Lisp Variables\]](#), page 44.
  - all Lisp objects that are created in C functions and that must be protected from freeing them. They are registered in the global list `gcprolist`. See [Section 14.3 \[GCPROing\]](#), page 86.
  - all local variables (i.e. their name fields `symbol` and old values `old_values`) that are bound during the evaluation by the Lisp engine. They are stored in `specbinding` structs pushed on a stack called `specpdl`. See [Section 17.2 \[Dynamic Binding; The specbinding Stack; Unwind-Protects\]](#), page 119.

- all catch blocks that the Lisp engine encounters during the evaluation cause the creation of structs `catchtag` inserted in the list `catchlist`. Their tag (`tag`) and value (`val` fields are freshly created objects and therefore have to be marked. See [Section 17.4 \[Catch and Throw\]](#), page 119.
  - every function application pushes new structs `backtrace` on the call stack of the Lisp engine (`backtrace_list`). The unique parts that have to be marked are the fields for each function (`function`) and all their arguments (`args`). See [Section 17.1 \[Evaluation\]](#), page 117.
  - all objects that are used by the redisplay engine that must not be freed are marked by a special function called `mark_redisplay` (in `redisplay.c`).
  - all objects created for profiling purposes are allocated by C functions instead of using the lisp allocation mechanisms. In order to receive the right ones during the sweep phase, they also have to be marked manually. That is done by the function `mark_profiling_info`
9. Hash tables in XEmacs belong to a kind of special objects that make use of a concept often called 'weak pointers'. To make a long story short, these kind of pointers are not followed during the estimation of the live objects during garbage collection. Any object referenced only by weak pointers is collected anyway, and the reference to it is cleared. In hash tables there are different usage patterns of them, manifesting in different types of hash tables, namely 'non-weak', 'weak', 'key-weak' and 'value-weak' (internally also 'key-car-weak' and 'value-car-weak') hash tables, each clearing entries depending on different conditions. More information can be found in the documentation to the function `make-hash-table`.

Because there are complicated dependency rules about when and what to mark while processing weak hash tables, the standard `marker` method is only active if it is marking non-weak hash tables. As soon as a weak component is in the table, the hash table entries are ignored while marking. Instead their marking is done each separately by the function `finish_marking_weak_hash_tables`. This function iterates over each hash table entry `hentries` for each weak hash table in `Vall_weak_hash_tables`. Depending on the type of a table, the appropriate action is performed. If a table is acting as `HASH_TABLE_KEY_WEAK`, and a key already marked, everything reachable from the `value` component is marked. If it is acting as a `HASH_TABLE_VALUE_WEAK` and the value component is already marked, the marking starts beginning only from the `key` component. If it is a `HASH_TABLE_KEY_CAR_WEAK` and the car of the key entry is already marked, we mark both the `key` and `value` components. Finally, if the table is of the type `HASH_TABLE_VALUE_CAR_WEAK` and the car of the value components is already marked, again both the `key` and the `value` components get marked.

Again, there are lists with comparable properties called weak lists. There exist different peculiarities of their types called `simple`, `assoc`, `key-assoc` and `value-assoc`. You can find further details about them in the description to the function `make-weak-list`. The scheme of their marking is similar: all weak lists are listed in `Qall_weak_lists`, therefore we iterate over them. The marking is advanced until we hit an already marked pair. Then we know that during a former run all the rest has been marked completely. Again, depending on the special type of the weak list, our jobs differ. If it is a `WEAK_LIST_SIMPLE` and the `elem` is marked, we mark the `cons` part. If it is a `WEAK_LIST_ASSOC` and not a pair or a pair with both marked car and cdr, we mark the

`cons` and the `elem`. If it is a `WEAK_LIST_KEY_ASSOC` and not a pair or a pair with a marked car of the `elem`, we mark the `cons` and the `elem`. Finally, if it is a `WEAK_LIST_VALUE_ASSOC` and not a pair or a pair with a marked cdr of the `elem`, we mark both the `cons` and the `elem`.

Since, by marking objects in reach from weak hash tables and weak lists, other objects could get marked, this perhaps implies further marking of other weak objects, both finishing functions are redone as long as yet unmarked objects get freshly marked.

10. After completing the special marking for the weak hash tables and for the weak lists, all entries that point to objects that are going to be swept in the further process are useless, and therefore have to be removed from the table or the list.

The function `prune_weak_hash_tables` does the job for weak hash tables. Totally unmarked hash tables are removed from the list `Vall_weak_hash_tables`. The other ones are treated more carefully by scanning over all entries and removing one as soon as one of the components `key` and `value` is unmarked.

The same idea applies to the weak lists. It is accomplished by `prune_weak_lists`: An unmarked list is pruned from `Vall_weak_lists` immediately. A marked list is treated more carefully by going over it and removing just the unmarked pairs.

11. The function `prune_specifiers` checks all listed specifiers held in `Vall_specifiers` and removes the ones from the lists that are unmarked.
12. All syntax tables are stored in a list called `Vall_syntax_tables`. The function `prune_syntax_tables` walks through it and unlinks the tables that are unmarked.
13. Next, we will attack the complete sweeping - the function `gc_sweep` which holds the predominance.
14. First, all the variables with respect to garbage collection are reset. `consing_since_gc` - the counter of the created cells since the last garbage collection - is set back to 0, and `gc_in_progress` is not `true` anymore.
15. In case the session is interactive, the displayed cursor and message are removed again.
16. The state of `gc_inhibit` is restored to the former value by unwinding the stack.
17. A small memory reserve is always held back that can be reached by `breathing_space`. If nothing more is left, we create a new reserve and exit.

### 14.4.3 mark\_object

The first thing that is checked while marking an object is whether the object is a real Lisp object `Lisp_Type_Record` or just an integer or a character. Integers and characters are the only two types that are stored directly - without another level of indirection, and therefore they don't have to be marked and collected. See [Chapter 8 \[How Lisp Objects Are Represented in C\]](#), page 29.

The second case is the one we have to handle. It is the one when we are dealing with a pointer to a Lisp object. But, there exist also three possibilities, that prevent us from doing anything while marking: The object is read only which prevents it from being garbage collected, i.e. marked (`C_READONLY_RECORD_HEADER`). The object in question is already marked, and need not be marked for the second time (checked by `MARKED_RECORD_HEADER_P`). If it is a special, unmarkable object (`UNMARKABLE_RECORD_HEADER_P`, apparently, these are objects that sit in some const space, and can therefore not be marked, see `this_one_is_unmarkable` in `alloc.c`).

Now, the actual marking is feasible. We do so by once using the macro `MARK_RECORD_HEADER` to mark the object itself (actually the special flag in the record header), and calling its special marker "method" `marker` if available. The marker method marks every other object that is in reach from our current object. Note, that these marker methods should not call `mark_object` recursively, but instead should return the next object from where further marking has to be performed.

In case another object was returned, as mentioned before, we reiterate the whole `mark_object` process beginning with this next object.

#### 14.4.4 `gc_sweep`

The job of this function is to free all unmarked records from memory. As we know, there are different types of objects implemented and managed, and consequently different ways to free them from memory. See [Section 14.1 \[Introduction to Allocation\], page 85](#).

We start with all objects stored through `lcrecords`. All bulkier objects are allocated and handled using that scheme of `lcrecords`. Each object is `malloced` separately instead of placing it in one of the contiguous frob blocks. All types that are currently stored using `lcrecords`'s `alloc_lcrecord` and `make_lcrecord_list` are the types: vectors, buffers, char-table, char-table-entry, console, weak-list, database, device, ldap, hash-table, command-builder, extent-auxiliary, extent-info, face, coding-system, frame, image-instance, glyph, popup-data, gui-item, keymap, charset, color\_instance, font\_instance, opaque, opaque-list, process, range-table, specifier, symbol-value-buffer-local, symbol-value-lisp-magic, symbol-value-varalias, toolbar-button, tootalk-message, tootalk-pattern, window, and window-configuration. We take care of them in the first place in order to be able to handle and to finalize items stored in them more easily. The function `sweep_lcrecords_1` as described below is doing the whole job for us. For a description about the internals: See [Section 14.7 \[lcrecords\], page 95](#).

Our next candidates are the other objects that behave quite differently than everything else: the strings. They consists of two parts, a fixed-size portion (`struct Lisp_String`) holding the string's length, its property list and a pointer to the second part, and the actual string data, which is stored in string-chars blocks comparable to frob blocks. In this block, the data is not only freed, but also a compression of holes is made, i.e. all strings are relocated together. See [Section 14.14 \[String\], page 100](#). This compacting phase is performed by the function `compact_string_chars`, the actual sweeping by the function `sweep_strings` is described below.

After that, the other types are swept step by step using functions `sweep_conses`, `sweep_bit_vectors_1`, `sweep_compiled_functions`, `sweep_floats`, `sweep_symbols`, `sweep_extents`, `sweep_markers` and `sweep_extents`. They are the fixed-size types conses, floats, compiled-functions, symbol, marker, extent, and event stored in so-called "frob blocks", and therefore we can basically do the same on every type objects, using the same macros, especially defined only to handle everything with respect to fixed-size blocks. The only fixed-size type that is not handled here are the fixed-size portion of strings, because we took special care of them earlier.

The only big exceptions are bit vectors stored differently and therefore treated differently by the function `sweep_bit_vectors_1` described later.

At first, we need some brief information about how these fixed-size types are managed in general, in order to understand how the sweeping is done. They have all a fixed size, and

are therefore stored in big blocks of memory - allocated at once - that can hold a certain amount of objects of one type. The macro `DECLARE_FIXED_TYPE_ALLOC` creates the suitable structures for every type. More precisely, we have the block struct (holding a pointer to the previous block `prev` and the objects in `block[]`), a pointer to current block (`current_..._block`) and its last index (`current_..._block_index`), and a pointer to the free list that will be created. Also a macro `FIXED_TYPE_FROM_BLOCK` plus some related macros exists that are used to obtain a new object, either from the free list `ALLOCATE_FIXED_TYPE_1` if there is an unused object of that type stored or by allocating a completely new block using `ALLOCATE_FIXED_TYPE_FROM_BLOCK`.

The rest works as follows: all of them define a macro `UNMARK_...` that is used to unmark the object. They define a macro `ADDITIONAL_FREE_...` that defines additional work that has to be done when converting an object from in use to not in use (so far, only markers use it in order to unchain them). Then, they all call the macro `SWEEP_FIXED_TYPE_BLOCK` instantiated with their type name and their struct name.

This call in particular does the following: we go over all blocks starting with the current moving towards the oldest. For each block, we look at every object in it. If the object already freed (checked with `FREE_STRUCT_P` using the first pointer of the object), or if it is set to read only (`C_READONLY_RECORD_HEADER_P`, nothing must be done. If it is unmarked (checked with `MARKED_RECORD_HEADER_P`), it is put in the free list and set free (using the macro `FREE_FIXED_TYPE`, otherwise it stays in the block, but is unmarked (by `UNMARK_...`). While going through one block, we note if the whole block is empty. If so, the whole block is freed (using `xfree`) and the free list state is set to the state it had before handling this block.

#### 14.4.5 `sweep_lcrecords_1`

After nullifying the complete lcrecord statistics, we go over all lcrecords two separate times. They are all chained together in a list with a head called `all_lcrecords`.

The first loop calls for each object its `finalizer` method, but only in the case that it is not read only (`C_READONLY_RECORD_HEADER_P`), it is not already marked (`MARKED_RECORD_HEADER_P`), it is not already in a free list (list of freed objects, field `free`) and finally it owns a finalizer method.

The second loop actually frees the appropriate objects again by iterating through the whole list. In case an object is read only or marked, it has to persist, otherwise it is manually freed by calling `xfree`. During this loop, the lcrecord statistics are kept up to date by calling `tick_lcrecord_stats` with the right arguments,

#### 14.4.6 `compact_string_chars`

The purpose of this function is to compact all the data parts of the strings that are held in so-called `string_chars_block`, i.e. the strings that do not exceed a certain maximal length.

The procedure with which this is done is as follows. We are keeping two positions in the `string_chars_blocks` using two pointer/integer pairs, namely `from_sb/from_pos` and `to_sb/to_pos`. They stand for the actual positions, from where to where, to copy the actually handled string.

While going over all chained `string_char_blocks` and their held strings, starting at `first_string_chars_block`, both pointers are advanced and eventually a string is copied from `from_sb` to `to_sb`, depending on the status of the pointed at strings.

More precisely, we can distinguish between the following actions.

- The string at `from_sb`'s position could be marked as free, which is indicated by an invalid pointer to the pointer that should point back to the fixed size string object, and which is checked by `FREE_STRUCT_P`. In this case, the `from_sb/from_pos` is advanced to the next string, and nothing has to be copied.
- Also, if a string object itself is unmarked, nothing has to be copied. We likewise advance the `from_sb/from_pos` pair as described above.
- In all other cases, we have a marked string at hand. The string data must be moved from the from-position to the to-position. In case there is not enough space in the actual `to_sb`-block, we advance this pointer to the beginning of the next block before copying. In case the from and to positions are different, we perform the actual copying using the library function `memmove`.

After compacting, the pointer to the current `string_chars_block`, sitting in `current_string_chars_block`, is reset on the last block to which we moved a string, i.e. `to_block`, and all remaining blocks (we know that they just carry garbage) are explicitly `x freed`.

#### 14.4.7 `sweep_strings`

The sweeping for the fixed sized string objects is essentially exactly the same as it is for all other fixed size types. As before, the freeing into the suitable free list is done by using the macro `SWEEP_FIXED_SIZE_BLOCK` after defining the right macros `UNMARK_string` and `ADDITIONAL_FREE_string`. These two definitions are a little bit special compared to the ones used for the other fixed size types.

`UNMARK_string` is defined the same way except some additional code used for updating the bookkeeping information.

For strings, `ADDITIONAL_FREE_string` has to do something in addition: in case, the string was not allocated in a `string_chars_block` because it exceeded the maximal length, and therefore it was `malloced` separately, we know also `x free` it explicitly.

#### 14.4.8 `sweep_bit_vectors_1`

Bit vectors are also one of the rare types that are `malloced` individually. Consequently, while sweeping, all further needless bit vectors must be freed by hand. This is done, as one might imagine, the expected way: since they are all registered in a list called `all_bit_vectors`, all elements of that list are traversed, all unmarked bit vectors are unlinked by calling `x free` and all of them become unmarked. In addition, the bookkeeping information used for garbage collector's output purposes is updated.

### 14.5 Integers and Characters

Integer and character Lisp objects are created from integers using the macros `XSETINT()` and `XSETCHAR()` or the equivalent functions `make_int()` and `make_char()`. (These are actually macros on most systems.) These functions basically just do some moving of bits around, since the integral value of the object is stored directly in the `Lisp_Object`.

`XSETINT()` and the like will truncate values given to them that are too big; i.e. you won't get the value you expected but the tag bits will at least be correct.

## 14.6 Allocation from Frob Blocks

The uninitialized memory required by a `Lisp_Object` of a particular type is allocated using `ALLOCATE_FIXED_TYPE()`. This only occurs inside of the lowest-level object-creating functions in 'alloc.c': `Fcons()`, `make_float()`, `Fmake_byte_code()`, `Fmake_symbol()`, `allocate_extent()`, `allocate_event()`, `Fmake_marker()`, and `make_uninit_string()`. The idea is that, for each type, there are a number of frob blocks (each 2K in size); each frob block is divided up into object-sized chunks. Each frob block will have some of these chunks that are currently assigned to objects, and perhaps some that are free. (If a frob block has nothing but free chunks, it is freed at the end of the garbage collection cycle.) The free chunks are stored in a free list, which is chained by storing a pointer in the first four bytes of the chunk. (Except for the free chunks at the end of the last frob block, which are handled using an index which points past the end of the last-allocated chunk in the last frob block.) `ALLOCATE_FIXED_TYPE()` first tries to retrieve a chunk from the free list; if that fails, it calls `ALLOCATE_FIXED_TYPE_FROM_BLOCK()`, which looks at the end of the last frob block for space, and creates a new frob block if there is none. (There are actually two versions of these macros, one of which is more defensive but less efficient and is used for error-checking.)

## 14.7 lrecords

[see 'lrecord.h']

All lrecords have at the beginning of their structure a `struct lrecord_header`. This just contains a type number and some flags, including the mark bit. All builtin type numbers are defined as constants in `enum lrecord_type`, to allow the compiler to generate more efficient code for type `P`. The type number, thru the `lrecord_implementation_table`, gives access to a `struct lrecord_implementation`, which is a structure containing method pointers and such. There is one of these for each type, and it is a global, constant, statically-declared structure that is declared in the `DEFINE_LRECORD_IMPLEMENTATION()` macro.

Simple lrecords (of type (b) above) just have a `struct lrecord_header` at their beginning. lrecords, however, actually have a `struct lcrecord_header`. This, in turn, has a `struct lrecord_header` at its beginning, so sanity is preserved; but it also has a pointer used to chain all lcrecords together, and a special ID field used to distinguish one lcrecord from another. (This field is used only for debugging and could be removed, but the space gain is not significant.)

Simple lrecords are created using `ALLOCATE_FIXED_TYPE()`, just like for other frob blocks. The only change is that the implementation pointer must be initialized correctly. (The implementation structure for an lrecord, or rather the pointer to it, is named `lrecord_float`, `lrecord_extent`, `lrecord_buffer`, etc.)

lcrecords are created using `alloc_lcrecord()`. This takes a size to allocate and an implementation pointer. (The size needs to be passed because some lcrecords, such as window configurations, are of variable size.) This basically just `malloc()`s the storage, initializes the `struct lcrecord_header`, and chains the lcrecord onto the head of the list of

all `lrecords`, which is stored in the variable `all_lrecords`. The calls to `alloc_lrecord()` generally occur in the lowest-level allocation function for each `lrecord` type.

Whenever you create an `lrecord`, you need to call either `DEFINE_LRECORD_IMPLEMENTATION()` or `DEFINE_LRECORD_SEQUENCE_IMPLEMENTATION()`. This needs to be specified in a `.c` file, at the top level. What this actually does is define and initialize the implementation structure for the `lrecord`. (And possibly declares a function `error_check_foo()` that implements the `XFOO()` macro when error-checking is enabled.) The arguments to the macros are the actual type name (this is used to construct the C variable name of the `lrecord` implementation structure and related structures using the `##` macro concatenation operator), a string that names the type on the Lisp level (this may not be the same as the C type name; typically, the C type name has underscores, while the Lisp string has dashes), various method pointers, and the name of the C structure that contains the object. The methods are used to encapsulate type-specific information about the object, such as how to print it or mark it for garbage collection, so that it's easy to add new object types without having to add a specific case for each new type in a bunch of different places.

The difference between `DEFINE_LRECORD_IMPLEMENTATION()` and `DEFINE_LRECORD_SEQUENCE_IMPLEMENTATION()` is that the former is used for fixed-size object types and the latter is for variable-size object types. Most object types are fixed-size; some complex types, however (e.g. window configurations), are variable-size. Variable-size object types have an extra method, which is called to determine the actual size of a particular object of that type. (Currently this is only used for keeping allocation statistics.)

For the purpose of keeping allocation statistics, the allocation engine keeps a list of all the different types that exist. Note that, since `DEFINE_LRECORD_IMPLEMENTATION()` is a macro that is specified at top-level, there is no way for it to initialize the global data structures containing type information, like `lrecord_implementations_table`. For this reason a call to `INIT_LRECORD_IMPLEMENTATION` must be added to the same source file containing `DEFINE_LRECORD_IMPLEMENTATION`, but instead of to the top level, to one of the init functions, typically `syms_of_foo.c`. `INIT_LRECORD_IMPLEMENTATION` must be called before an object of this type is used.

The type number is also used to index into an array holding the number of objects of each type and the total memory allocated for objects of that type. The statistics in this array are computed during the sweep stage. These statistics are returned by the call to `garbage-collect`.

Note that for every type defined with a `DEFINE_LRECORD_*` macro, there needs to be a `DECLARE_LRECORD_IMPLEMENTATION()` somewhere in a `.h` file, and this `.h` file needs to be included by `inline.c`.

Furthermore, there should generally be a set of `XFOOBAR()`, `FOOBARP()`, etc. macros in a `.h` (or occasionally `.c`) file. To create one of these, copy an existing model and modify as necessary.

**Please note:** If you define an `lrecord` in an external dynamically-loaded module, you must use `DECLARE_EXTERNAL_LRECORD`, `DEFINE_EXTERNAL_LRECORD_IMPLEMENTATION`, and `DEFINE_EXTERNAL_LRECORD_SEQUENCE_IMPLEMENTATION` instead of the non-EXTERNAL forms. These macros will dynamically add new type numbers to the global enum that

records them, whereas the non-EXTERNAL forms assume that the programmer has already inserted the correct type numbers into the enum's code at compile-time.

The various methods in the lrecord implementation structure are:

1. A **mark** method. This is called during the marking stage and passed a function pointer (usually the `mark_object()` function), which is used to mark an object. All Lisp objects that are contained within the object need to be marked by applying this function to them. The mark method should also return a Lisp object, which should be either `nil` or an object to mark. (This can be used in lieu of calling `mark_object()` on the object, to reduce the recursion depth, and consequently should be the most heavily nested sub-object, such as a long list.)

**Please note:** When the mark method is called, garbage collection is in progress, and special precautions need to be taken when accessing objects; see section (B) above.

If your mark method does not need to do anything, it can be `NULL`.

2. A **print** method. This is called to create a printed representation of the object, whenever `princ`, `prin1`, or the like is called. It is passed the object, a stream to which the output is to be directed, and an `escapeflag` which indicates whether the object's printed representation should be **escaped** so that it is readable. (This corresponds to the difference between `princ` and `prin1`.) Basically, **escaped** means that strings will have quotes around them and confusing characters in the strings such as quotes, backslashes, and newlines will be backslashed; and that special care will be taken to make symbols print in a readable fashion (e.g. symbols that look like numbers will be backslashed). Other readable objects should perhaps pass `escapeflag` on when sub-objects are printed, so that readability is preserved when necessary (or if not, always pass in a 1 for `escapeflag`). Non-readable objects should in general ignore `escapeflag`, except that some use it as an indication that more verbose output should be given.

Sub-objects are printed using `print_internal()`, which takes exactly the same arguments as are passed to the print method.

Literal C strings should be printed using `write_c_string()`, or `write_string_1()` for non-null-terminated strings.

Functions that do not have a readable representation should check the `print_readably` flag and signal an error if it is set.

If you specify `NULL` for the print method, the `default_object_printer()` will be used.

3. A **nalize** method. This is called at the beginning of the sweep stage on lrecords that are about to be freed, and should be used to perform any extra object cleanup. This typically involves freeing any extra `malloc()`ed memory associated with the object, releasing any operating-system and window-system resources associated with the object (e.g. pixmaps, fonts), etc.

The finalize method can be `NULL` if nothing needs to be done.

**WARNING #1:** The finalize method is also called at the end of the dump phase; this time with the `for_disksave` parameter set to non-zero. The object is *not* about to disappear, so you have to make sure to *not* free any extra `malloc()`ed memory if you're going to need it later. (Also, signal an error if there are any operating-system and window-system resources here, because they can't be dumped.)

Finalize methods should, as a rule, set to zero any pointers after they've been freed, and check to make sure pointers are not zero before freeing. Although I'm pretty sure that finalize methods are not called twice on the same object (except for the `for_disksave` proviso), we've gotten nastily burned in some cases by not doing this.

WARNING #2: The finalize method is *only* called for lrecords, *not* for simply lrecords. If you need a finalize method for simple lrecords, you have to stick it in the `ADDITIONAL_FREE_foo()` macro in `'alloc.c'`.

WARNING #3: Things are in an *extremely* bizarre state when `ADDITIONAL_FREE_foo()` is called, so you have to be incredibly careful when writing one of these functions. See the comment in `gc_sweep()`. If you ever have to add one of these, consider using an lrecord or dealing with the problem in a different fashion.

4. An `equal` method. This compares the two objects for similarity, when `equal` is called. It should compare the contents of the objects in some reasonable fashion. It is passed the two objects and a `depth` value, which is used to catch circular objects. To compare sub-Lisp-objects, call `internal_equal()` and bump the depth value by one. If this value gets too high, a `circular-object` error will be signaled.

If this is NULL, objects are `equal` only when they are `eq`, i.e. identical.

5. A `hash` method. This is used to hash objects when they are to be compared with `equal`. The rule here is that if two objects are `equal`, they *must* hash to the same value; i.e. your hash function should use some subset of the sub-fields of the object that are compared in the "equal" method. If you specify this method as NULL, the object's pointer will be used as the hash, which will *fail* if the object has an `equal` method, so don't do this.

To hash a sub-Lisp-object, call `internal_hash()`. Bump the depth by one, just like in the "equal" method.

To convert a Lisp object directly into a hash value (using its pointer), use `LISP_HASH()`. This is what happens when the hash method is NULL.

To hash two or more values together into a single value, use `HASH2()`, `HASH3()`, `HASH4()`, etc.

6. `getprop`, `putprop`, `remprop`, and `plist` methods. These are used for object types that have properties. I don't feel like documenting them here. If you create one of these objects, you have to use different macros to define them, i.e. `DEFINE_LRECORD_IMPLEMENTATION_WITH_PROPS()` or `DEFINE_LRECORD_SEQUENCE_IMPLEMENTATION_WITH_PROPS()`.
7. A `size.in.bytes` method, when the object is of variable-size. (i.e. declared with a `_SEQUENCE_IMPLEMENTATION` macro.) This should simply return the object's size in bytes, exactly as you might expect. For an example, see the methods for window configurations and opaques.

## 14.8 Low-level allocation

Memory that you want to allocate directly should be allocated using `xmalloc()` rather than `malloc()`. This implements error-checking on the return value, and once upon a time did some more vital stuff (i.e. `BLOCK_INPUT`, which is no longer necessary). Free using `xfree()`, and realloc using `xrealloc()`. Note that `xmalloc()` will do a non-local exit if

the memory can't be allocated. (Many functions, however, do not expect this, and thus XEmacs will likely crash if this happens. **This is a bug.** If you can, you should strive to make your function handle this OK. However, it's difficult in the general circumstance, perhaps requiring extra `unwind-protects` and such.)

Note that XEmacs provides two separate replacements for the standard `malloc()` library function. These are called **old GNU malloc** (`'malloc.c'`) and **new GNU malloc** (`'gmalloc.c'`), respectively. New GNU malloc is better in pretty much every way than old GNU malloc, and should be used if possible. (It used to be that on some systems, the old one worked but the new one didn't. I think this was due specifically to a bug in SunOS, which the new one now works around; so I don't think the old one ever has to be used any more.) The primary difference between both of these mallocs and the standard system malloc is that they are much faster, at the expense of increased space. The basic idea is that memory is allocated in fixed chunks of powers of two. This allows for basically constant malloc time, since the various chunks can just be kept on a number of free lists. (The standard system malloc typically allocates arbitrary-sized chunks and has to spend some time, sometimes a significant amount of time, walking the heap looking for a free block to use and cleaning things up.) The new GNU malloc improves on things by allocating large objects in chunks of 4096 bytes rather than in ever larger powers of two, which results in ever larger wastage. There is a slight speed loss here, but it's of doubtful significance.

NOTE: Apparently there is a third-generation GNU malloc that is significantly better than the new GNU malloc, and should probably be included in XEmacs.

There is also the relocating allocator, `'ralloc.c'`. This actually moves blocks of memory around so that the `sbrk()` pointer shrunk and virtual memory released back to the system. On some systems, this is a big win. On all systems, it causes a noticeable (and sometimes huge) speed penalty, so I turn it off by default. `'ralloc.c'` only works with the new GNU malloc in `'gmalloc.c'`. There are also two versions of `'ralloc.c'`, one that uses `mmap()` rather than block copies to move data around. This purports to be faster, although that depends on the amount of data that would have had to be block copied and the system-call overhead for `mmap()`. I don't know exactly how this works, except that the relocating-allocation routines are pretty much used only for the memory allocated for a buffer, which is the biggest consumer of space, esp. of space that may get freed later.

Note that the GNU mallocs have some "memory warning" facilities. XEmacs taps into them and issues a warning through the standard warning system, when memory gets to 75%, 85%, and 95% full. (On some systems, the memory warnings are not functional.)

Allocated memory that is going to be used to make a Lisp object is created using `allocate_lisp_storage()`. This just calls `xmalloc()`. It used to verify that the pointer to the memory can fit into a Lisp word, before the current Lisp object representation was introduced. `allocate_lisp_storage()` is called by `alloc_lcrecord()`, `ALLOCATE_FIXED_TYPE()`, and the vector and bit-vector creation routines. These routines also call `INCREMENT_CONS_COUNTER()` at the appropriate times; this keeps statistics on how much memory is allocated, so that garbage-collection can be invoked when the threshold is reached.

## 14.9 Cons

Conses are allocated in standard frob blocks. The only thing to note is that conses can be explicitly freed using `free_cons()` and associated functions `free_list()` and `free_`

`alist()`. This immediately puts the conses onto the cons free list, and decrements the statistics on memory allocation appropriately. This is used to good effect by some extremely commonly-used code, to avoid generating extra objects and thereby triggering GC sooner. However, you have to be *extremely* careful when doing this. If you mess this up, you will get BADLY BURNED, and it has happened before.

## 14.10 Vector

As mentioned above, each vector is `malloc()`ed individually, and all are threaded through the variable `all_vectors`. Vectors are marked strangely during garbage collection, by kludging the size field. Note that the `struct Lisp_Vector` is declared with its `contents` field being a *stretchy* array of one element. It is actually `malloc()`ed with the right size, however, and access to any element through the `contents` array works fine.

## 14.11 Bit Vector

Bit vectors work exactly like vectors, except for more complicated code to access an individual bit, and except for the fact that bit vectors are lrecords while vectors are not. (The only difference here is that there's an lrecord implementation pointer at the beginning and the tag field in bit vector Lisp words is "lrecord" rather than "vector".)

## 14.12 Symbol

Symbols are also allocated in frob blocks. Symbols in the awful horrible obarray structure are chained through their `next` field.

Remember that `intern` looks up a symbol in an obarray, creating one if necessary.

## 14.13 Marker

Markers are allocated in frob blocks, as usual. They are kept in a buffer unordered, but in a doubly-linked list so that they can easily be removed. (Formerly this was a singly-linked list, but in some cases garbage collection took an extraordinarily long time due to the  $O(N^2)$  time required to remove lots of markers from a buffer.) Markers are removed from a buffer in the finalize stage, in `ADDITIONAL_FREE_marker()`.

## 14.14 String

As mentioned above, strings are a special case. A string is logically two parts, a fixed-size object (containing the length, property list, and a pointer to the actual data), and the actual data in the string. The fixed-size object is a `struct Lisp_String` and is allocated in frob blocks, as usual. The actual data is stored in special **string-chars blocks** which are 8K blocks of memory. Currently-allocated strings are simply laid end to end in these string-chars blocks, with a pointer back to the `struct Lisp_String` stored before each string in the string-chars block. When a new string needs to be allocated, the remaining space at the end of the last string-chars block is used if there's enough, and a new string-chars block is created otherwise.

There are never any holes in the string-chars blocks due to the string compaction and relocation that happens at the end of garbage collection. During the sweep stage of garbage collection, when objects are reclaimed, the garbage collector goes through all string-chars

blocks, looking for unused strings. Each chunk of string data is preceded by a pointer to the corresponding `struct Lisp_String`, which indicates both whether the string is used and how big the string is, i.e. how to get to the next chunk of string data. Holes are compressed by block-copying the next string into the empty space and relocating the pointer stored in the corresponding `struct Lisp_String`. **This means you have to be careful with strings in your code.** See the section above on GCPR0ing.

Note that there is one situation not handled: a string that is too big to fit into a string-chars block. Such strings, called **big strings**, are all `malloc()`ed as their own block. (##### Although it would make more sense for the threshold for big strings to be somewhat lower, e.g. 1/2 or 1/4 the size of a string-chars block. It seems that this was indeed the case formerly—indeed, the threshold was set at 1/8—but Mly forgot about this when rewriting things for 19.8.)

Note also that the string data in string-chars blocks is padded as necessary so that proper alignment constraints on the `struct Lisp_String` back pointers are maintained.

Finally, strings can be resized. This happens in Mule when a character is substituted with a different-length character, or during modeline frobbing. (You could also export this to Lisp, but it's not done so currently.) Resizing a string is a potentially tricky process. If the change is small enough that the padding can absorb it, nothing other than a simple memory move needs to be done. Keep in mind, however, that the string can't shrink too much because the offset to the next string in the string-chars block is computed by looking at the length and rounding to the nearest multiple of four or eight. If the string would shrink or expand beyond the correct padding, new string data needs to be allocated at the end of the last string-chars block and the data moved appropriately. This leaves some dead string data, which is marked by putting a special marker of `0xFFFFFFFF` in the `struct Lisp_String` pointer before the data (there's no real `struct Lisp_String` to point to and relocate), and storing the size of the dead string data (which would normally be obtained from the now-non-existent `struct Lisp_String`) at the beginning of the dead string data gap. The string compactor recognizes this special `0xFFFFFFFF` marker and handles it correctly.

## 14.15 Compiled Function

Not yet documented.



## 15 Dumping

### 15.1 What is dumping and its justification

The C code of XEmacs is just a Lisp engine with a lot of built-in primitives useful for writing an editor. The editor itself is written mostly in Lisp, and represents around 100K lines of code. Loading and executing the initialization of all this code takes a bit a time (five to ten times the usual startup time of current `xemacs`) and requires having all the lisp source files around. Having to reload them each time the editor is started would not be acceptable.

The traditional solution to this problem is called dumping: the build process first creates the lisp engine under the name `'temacs'`, then runs it until it has finished loading and initializing all the lisp code, and eventually creates a new executable called `'xemacs'` including both the object code in `'temacs'` and all the contents of the memory after the initialization.

This solution, while working, has a huge problem: the creation of the new executable from the actual contents of memory is an extremely system-specific process, quite error-prone, and which interferes with a lot of system libraries (like `malloc`). It is even getting worse nowadays with libraries using constructors which are automatically called when the program is started (even before `main()`) which tend to crash when they are called multiple times, once before dumping and once after (IRIX 6.x `libz.so` pulls in some C++ image libraries thru dependencies which have this problem). Writing the dumper is also one of the most difficult parts of porting XEmacs to a new operating system. Basically, 'dumping' is an operation that is just not officially supported on many operating systems.

The aim of the portable dumper is to solve the same problem as the system-specific dumper, that is to be able to reload quickly, using only a small number of files, the fully initialized lisp part of the editor, without any system-specific hacks.

### 15.2 Overview

The portable dumping system has to:

1. At dump time, write all initialized, non-quickly-rebuildable data to a file [Note: currently named `'xemacs.dmp'`, but the name will change], along with all informations needed for the reloading.
2. When starting `xemacs`, reload the dump file, relocate it to its new starting address if needed, and reinitialize all pointers to this data. Also, rebuild all the quickly rebuildable data.

### 15.3 Data descriptions

The more complex task of the dumper is to be able to write lisp objects (`lrecords`) and C structs to disk and reload them at a different address, updating all the pointers they include in the process. This is done by using external data descriptions that give information about the layout of the structures in memory.

The specification of these descriptions is in `lrecord.h`. A description of an `lrecord` is an array of `struct lrecord_description`. Each of these structs include a type, an offset in the structure and some optional parameters depending on the type. For instance, here is the string description:

```

static const struct lrecord_description string_description[] = {
  { XD_BYTECOUNT,      offsetof (Lisp_String, size) },
  { XD_OPAQUE_DATA_PTR,  offsetof (Lisp_String, data), XD_INDIRECT(0, 1) },
  { XD_LISP_OBJECT,     offsetof (Lisp_String, plist) },
  { XD_END }
};

```

The first line indicates a member of type `Bytecount`, which is used by the next, indirect directive. The second means "there is a pointer to some opaque data in the field `data`". The length of said data is given by the expression `XD_INDIRECT(0, 1)`, which means "the value in the 0th line of the description (welcome to C) plus one". The third line means "there is a `Lisp-Object` member `plist` in the `Lisp-String` structure". `XD_END` then ends the description.

This gives us all the information we need to move around what is pointed to by a structure (C or `lrecord`) and, by transitivity, everything that it points to. The only missing information for dumping is the size of the structure. For `lrecords`, this is part of the `lrecord_implementation`, so we don't need to duplicate it. For C structures we use a `struct_description`, which includes a size field and a pointer to an associated array of `lrecord_description`.

## 15.4 Dumping phase

Dumping is done by calling the function `pdump()` (in `dumper.c`) which is invoked from `Fdump_emacs` (in `emacs.c`). This function performs a number of tasks.

### 15.4.1 Object inventory

The first task is to build the list of the objects to dump. This includes:

- lisp objects
- C structures

We end up with one `pdump_entry_list_elmt` per object group (arrays of C structs are kept together) which includes a pointer to the first object of the group, the per-object size and the count of objects in the group, along with some other information which is initialized later.

These entries are linked together in `pdump_entry_list` structures and can be enumerated thru either:

1. the `pdump_object_table`, an array of `pdump_entry_list`, one per `lrecord` type, indexed by type number.
2. the `pdump_opaque_data_list`, used for the opaque data which does not include pointers, and hence does not need descriptions.
3. the `pdump_struct_table`, which is a vector of `struct_description/pdump_entry_list` pairs, used for non-opaque C structures.

This uses a marking strategy similar to the garbage collector. Some differences though:

1. We do not use the mark bit (which does not exist for C structures anyway); we use a big hash table instead.

2. We do not use the mark function of lrecords but instead rely on the external descriptions. This happens essentially because we need to follow pointers to C structures and opaque data in addition to Lisp\_Object members.

This is done by `pdump_register_object()`, which handles Lisp\_Object variables, and `pdump_register_struct()` which handles C structures, which both delegate the description management to `pdump_register_sub()`.

The hash table doubles as a map object to `pdump_entry_list_elt` (i.e. allows us to look up a `pdump_entry_list_elt` with the object it points to). Entries are added with `pdump_add_entry()` and looked up with `pdump_get_entry()`. There is no need for entry removal. The hash value is computed quite simply from the object pointer by `pdump_make_hash()`.

The roots for the marking are:

1. the `staticpro`'ed variables (there is a special `staticpro_nodump()` call for protected variables we do not want to dump).
2. the variables registered via `dump_add_root_object` (`staticpro()` is equivalent to `staticpro_nodump() + dump_add_root_object()`).
3. the variables registered via `dump_add_root_struct_ptr`, each of which points to a C structure.

This does not include the GCPRO'ed variables, the `specbinds`, the `catchtags`, the `backlist`, the `redisplay` or the `profiling` info, since we do not want to rebuild the actual chain of lisp calls which end up to the `dump-emacs` call, only the global variables.

Weak lists and weak hash tables are dumped as if they were their non-weak equivalent (without changing their type, of course). This has not yet been a problem.

## 15.4.2 Address allocation

The next step is to allocate the offsets of each of the objects in the final dump file. This is done by `pdump_allocate_offset()` which is called indirectly by `pdump_scan_by_alignment()`.

The strategy to deal with alignment problems uses these facts:

1. real world alignment requirements are powers of two.
2. the C compiler is required to adjust the size of a struct so that you can have an array of them next to each other. This means you can have an upper bound of the alignment requirements of a given structure by looking at which power of two its size is a multiple.
3. the non-variant part of variable size lrecords has an alignment requirement of 4.

Hence, for each lrecord type, C struct type or opaque data block the alignment requirement is computed as a power of two, with a minimum of  $2^2$  for lrecords. `pdump_scan_by_alignment()` then scans all the `pdump_entry_list_elt`'s, the ones with the highest requirements first. This ensures the best packing.

The maximum alignment requirement we take into account is  $2^8$ .

`pdump_allocate_offset()` only has to do a linear allocation, starting at offset 256 (this leaves room for the header and keeps the alignments happy).

### 15.4.3 The header

The next step creates the file and writes a header with a signature and some random information in it. The `reloc_address` field, which indicates at which address the file should be loaded if we want to avoid post-reload relocation, is set to 0. It then seeks to offset 256 (base offset for the objects).

### 15.4.4 Data dumping

The data is dumped in the same order as the addresses were allocated by `pdump_dump_data()`, called from `pdump_scan_by_alignment()`. This function copies the data to a temporary buffer, relocates all pointers in the object to the addresses allocated in step Address Allocation, and writes it to the file. Using the same order means that, if we are careful with lrecords whose size is not a multiple of 4, we are ensured that the object is always written at the offset in the file allocated in step Address Allocation.

### 15.4.5 Pointers dumping

A bunch of tables needed to reassign properly the global pointers are then written. They are:

1. the `pdump_root_struct_ptrs` dynarr
2. the `pdump_opaques` dynarr
3. a vector of all the offsets to the objects in the file that include a description (for faster relocation at reload time)
4. the `pdump_root_objects` and `pdump_weak_object_chains` dynarrs.

For each of the dynarrs we write both the pointer to the variables and the relocated offset of the object they point to. Since these variables are global, the pointers are still valid when restarting the program and are used to regenerate the global pointers.

The `pdump_weak_object_chains` dynarr is a special case. The variables it points to are the head of weak linked lists of lisp objects of the same type. Not all objects of this list are dumped so the relocated pointer we associate with them points to the first dumped object of the list, or `Qnil` if none is available. This is also the reason why they are not used as roots for the purpose of object enumeration.

Some very important information like the `staticpros` and `lrecord_implementations_table` are handled indirectly using `dump_add_opaque` or `dump_add_root_struct_ptr`.

This is the end of the dumping part.

## 15.5 Reloading phase

### 15.5.1 File loading

The file is `mmap`'ed in memory (which ensures a `PAGESIZE` alignment, at least 4096), or if `mmap` is unavailable or fails, a 256-bytes aligned `malloc` is done and the file is loaded.

Some variables are reinitialized from the values found in the header.

The difference between the actual loading address and the `reloc_address` is computed and will be used for all the relocations.

### 15.5.2 Putting back the pdump \_opagues

The memory contents are restored in the obvious and trivial way.

### 15.5.3 Putting back the pdump \_root \_struct \_ptrs

The variables pointed to by pdump\_root\_struct\_ptrs in the dump phase are reset to the right relocated object addresses.

### 15.5.4 Object relocation

All the objects are relocated using their description and their offset by pdump\_reloc\_one. This step is unnecessary if the reloc\_address is equal to the file loading address.

### 15.5.5 Putting back the pdump \_root \_objects and pdump \_weak \_object \_chains

Same as Putting back the pdump\_root\_struct\_ptrs.

### 15.5.6 Reorganize the hash tables

Since some of the hash values in the lisp hash tables are address-dependent, their layout is now wrong. So we go through each of them and have them resorted by calling pdump\_reorganize\_hash\_table.

## 15.6 Remaining issues

The build process will have to start a post-dump xemacs, ask it the loading address (which will, hopefully, be always the same between different xemacs invocations) and relocate the file to the new address. This way the object relocation phase will not have to be done, which means no writes in the objects and that, because of the use of mmap, the dumped data will be shared between all the xemacs running on the computer.

Some executable signature will be necessary to ensure that a given dump file is really associated with a given executable, or random crashes will occur. Maybe a random number set at compile or configure time thru a define. This will also allow for having differently-compiled xemacsen on the same system (mule and no-mule comes to mind).

The DOC file contents should probably end up in the dump file.



## 16 Events and the Event Loop

### 16.1 Introduction to Events

An event is an object that encapsulates information about an interesting occurrence in the operating system. Events are generated either by user action, direct (e.g. typing on the keyboard or moving the mouse) or indirect (moving another window, thereby generating an expose event on an Emacs frame), or as a result of some other typically asynchronous action happening, such as output from a subprocess being ready or a timer expiring. Events come into the system in an asynchronous fashion (typically through a callback being called) and are converted into a synchronous event queue (first-in, first-out) in a process that we will call **collection**.

Note that each application has its own event queue. (It is immaterial whether the collection process directly puts the events in the proper application's queue, or puts them into a single system queue, which is later split up.)

The most basic level of event collection is done by the operating system or window system. Typically, XEmacs does its own event collection as well. Often there are multiple layers of collection in XEmacs, with events from various sources being collected into a queue, which is then combined with other sources to go into another queue (i.e. a second level of collection), with perhaps another level on top of this, etc.

XEmacs has its own types of events (called **Emacs events**), which provides an abstract layer on top of the system-dependent nature of the most basic events that are received. Part of the complex nature of the XEmacs event collection process involves converting from the operating-system events into the proper Emacs events—there may not be a one-to-one correspondence.

Emacs events are documented in `'events.h'`; I'll discuss them later.

### 16.2 Main Loop

The **command loop** is the top-level loop that the editor is always running. It loops endlessly, calling `next-event` to retrieve an event and `dispatch-event` to execute it. `dispatch-event` does the appropriate thing with non-user events (process, timeout, magic, eval, mouse motion); this involves calling a Lisp handler function, redrawing a newly-exposed part of a frame, reading subprocess output, etc. For user events, `dispatch-event` looks up the event in relevant keymaps or menubars; when a full key sequence or menubar selection is reached, the appropriate function is executed. `dispatch-event` may have to keep state across calls; this is done in the “command-builder” structure associated with each console (remember, there's usually only one console), and the engine that looks up keystrokes and constructs full key sequences is called the **command builder**. This is documented elsewhere.

The guts of the command loop are in `command_loop_1()`. This function doesn't catch errors, though—that's the job of `command_loop_2()`, which is a condition-case (i.e. error-trapping) wrapper around `command_loop_1()`. `command_loop_1()` never returns, but may get thrown out of.

When an error occurs, `cmd_error()` is called, which usually invokes the Lisp error handler in `command-error`; however, a default error handler is provided if `command-error`

is `nil` (e.g. during startup). The purpose of the error handler is simply to display the error message and do associated cleanup; it does not need to throw anywhere. When the error handler finishes, the condition-case in `command_loop_2()` will finish and `command_loop_2()` will reinvoke `command_loop_1()`.

`command_loop_2()` is invoked from three places: from `initial_command_loop()` (called from `main()` at the end of internal initialization), from the Lisp function `recursive-edit`, and from `call_command_loop()`.

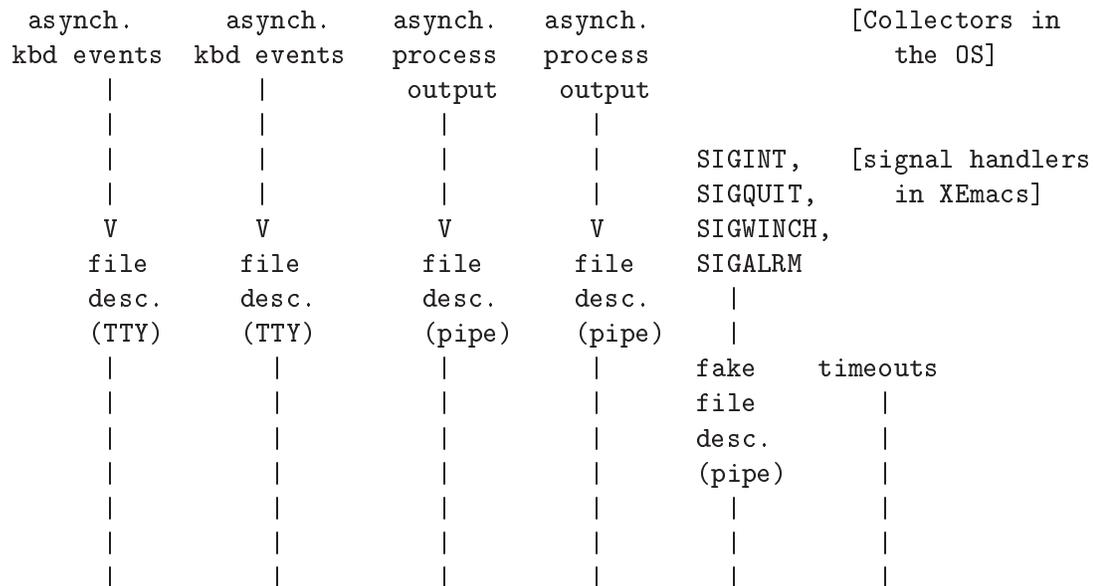
`call_command_loop()` is called when a macro is started and when the minibuffer is entered; normal termination of the macro or minibuffer causes a throw out of the recursive command loop. (To `execute-kbd-macro` for macros and `exit` for minibuffers. Note also that the low-level minibuffer-entering function, `read-minibuffer-internal`, provides its own error handling and does not need `command_loop_2()`'s error encapsulation; so it tells `call_command_loop()` to invoke `command_loop_1()` directly.)

Note that both `read-minibuffer-internal` and `recursive-edit` set up a catch for `exit`; this is why `abort-recursive-edit`, which throws to this catch, exits out of either one.

`initial_command_loop()`, called from `main()`, sets up a catch for `top-level` when invoking `command_loop_2()`, allowing functions to throw all the way to the top level if they really need to. Before invoking `command_loop_2()`, `initial_command_loop()` calls `top_level_1()`, which handles all of the startup stuff (creating the initial frame, handling the command-line options, loading the user's `.emacs` file, etc.). The function that actually does this is in Lisp and is pointed to by the variable `top-level`; normally this function is `normal-top-level`. `top_level_1()` is just an error-handling wrapper similar to `command_loop_2()`. Note also that `initial_command_loop()` sets up a catch for `top-level` when invoking `top_level_1()`, just like when it invokes `command_loop_2()`.

## 16.3 Specifics of the Event Gathering Mechanism

Here is an approximate diagram of the collection processes at work in XEmacs, under TTY's (TTY's are simpler than X so we'll look at this first):

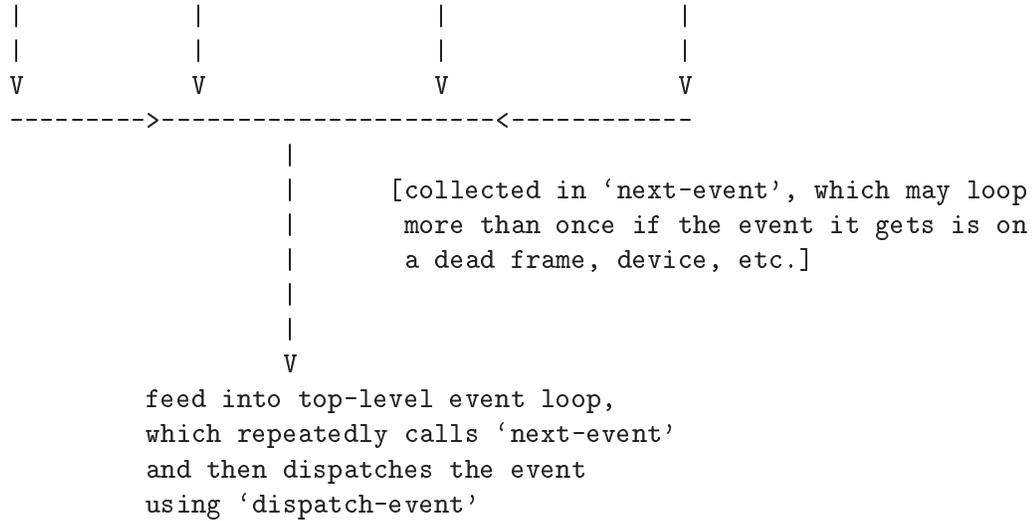












## 16.4 Specifics About the Emacs Event

## 16.5 The Event Stream Callback Routines

## 16.6 Other Event Loop Functions

`detect_input_pending()` and `input-pending-p` look for input by calling `event_stream->event_pending_p` and looking in `[V]unread-command-event` and the `command_event_queue` (they do not check for an executing keyboard macro, though).

`discard-input` cancels any command events pending (and any keyboard macros currently executing), and puts the others onto the `command_event_queue`. There is a comment about a “race condition”, which is not a good sign.

`next-command-event` and `read-char` are higher-level interfaces to `next-event`. `next-command-event` gets the next `commandevent` (i.e. keypress, mouse event, menu selection, or scrollbar action), calling `dispatch-event` on any others. `read-char` calls `next-command-event` and uses `event_to_character()` to return the character equivalent. With the right kind of input method support, it is possible for (read-char) to return a Kanji character.

## 16.7 Converting Events

`character_to_event()`, `event_to_character()`, `event-to-character`, and `character-to-event` convert between characters and keypress events corresponding to the characters. If the event was not a keypress, `event_to_character()` returns `-1` and `event-to-character` returns `nil`. These functions convert between character representation and the split-up event representation (keysym plus mod keys).

## 16.8 Dispatching Events; The Command Builder

Not yet documented.



## 17 Evaluation; Stack Frames; Bindings

### 17.1 Evaluation

`Feval()` evaluates the form (a Lisp object) that is passed to it. Note that evaluation is only non-trivial for two types of objects: symbols and conses. A symbol is evaluated simply by calling `symbol-value` on it and returning the value.

Evaluating a cons means calling a function. First, `eval` checks to see if garbage-collection is necessary, and calls `garbage_collect_1()` if so. It then increases the evaluation depth by 1 (`lisp_eval_depth`, which is always less than `max_lisp_eval_depth`) and adds an element to the linked list of `struct backtrace`'s (`backtrace_list`). Each such structure contains a pointer to the function being called plus a list of the function's arguments. Originally these values are stored unevalled, and as they are evaluated, the backtrace structure is updated. Garbage collection pays attention to the objects pointed to in the backtrace structures (garbage collection might happen while a function is being called or while an argument is being evaluated, and there could easily be no other references to the arguments in the argument list; once an argument is evaluated, however, the unevalled version is not needed by `eval`, and so the backtrace structure is changed).

At this point, the function to be called is determined by looking at the car of the cons (if this is a symbol, its function definition is retrieved and the process repeated). The function should then consist of either a `Lisp_Subr` (built-in function written in C), a `Lisp_Compiled_Function` object, or a cons whose car is one of the symbols `autoload`, `macro` or `lambda`.

If the function is a `Lisp_Subr`, the lisp object points to a `struct Lisp_Subr` (created by `DEFUN()`), which contains a pointer to the C function, a minimum and maximum number of arguments (or possibly the special constants `MANY` or `UNEVALLED`), a pointer to the symbol referring to that subr, and a couple of other things. If the subr wants its arguments `UNEVALLED`, they are passed raw as a list. Otherwise, an array of evaluated arguments is created and put into the backtrace structure, and either passed whole (`MANY`) or each argument is passed as a C argument.

If the function is a `Lisp_Compiled_Function`, `funcall_compiled_function()` is called. If the function is a lambda list, `funcall_lambda()` is called. If the function is a macro, [..... fill in] is done. If the function is an `autoload`, `do_autoload()` is called to load the definition and then `eval` starts over [explain this more].

When `Feval()` exits, the evaluation depth is reduced by one, the debugger is called if appropriate, and the current backtrace structure is removed from the list.

Both `funcall_compiled_function()` and `funcall_lambda()` need to go through the list of formal parameters to the function and bind them to the actual arguments, checking for `&rest` and `&optional` symbols in the formal parameters and making sure the number of actual arguments is correct. `funcall_compiled_function()` can do this a little more efficiently, since the formal parameter list can be checked for sanity when the compiled function object is created.

`funcall_lambda()` simply calls `Fprogn` to execute the code in the lambda list.

`funcall_compiled_function()` calls the real byte-code interpreter `execute_optimized_program()` on the byte-code instructions, which are converted into an internal form for faster execution.

When a compiled function is executed for the first time by `funcall_compiled_function()`, or during the dump phase of building XEmacs, the byte-code instructions are converted from a `Lisp_String` (which is inefficient to access, especially in the presence of MULE) into a `Lisp_Opaque` object containing an array of unsigned char, which can be directly executed by the byte-code interpreter. At this time the byte code is also analyzed for validity and transformed into a more optimized form, so that `execute_optimized_program()` can really fly.

Here are some of the optimizations performed by the internal byte-code transformer:

1. References to the `constants` array are checked for out-of-range indices, so that the byte interpreter doesn't have to.
2. References to the `constants` array that will be used as a Lisp variable are checked for being correct non-constant (i.e. not `t`, `nil`, or `keywordp`) symbols, so that the byte interpreter doesn't have to.
3. The maximum number of variable bindings in the byte-code is pre-computed, so that space on the `specpdl` stack can be pre-reserved once for the whole function execution.
4. All byte-code jumps are relative to the current program counter instead of the start of the program, thereby saving a register.
5. One-byte relative jumps are converted from the byte-code form of unsigned chars offset by 127 to machine-friendly signed chars.

Of course, this transformation of the `instructions` should not be visible to the user, so `Fcompiled_function_instructions()` needs to know how to convert the optimized opaque object back into a Lisp string that is identical to the original string from the `.elc` file. (Actually, the resulting string may (rarely) contain slightly different, yet equivalent, byte code.)

`Ffuncall()` implements Lisp `funcall`. (`funcall fun x1 x2 x3 ...`) is equivalent to `(eval (list fun (quote x1) (quote x2) (quote x3) ...))`. `Ffuncall()` contains its own code to do the evaluation, however, and is very similar to `Feval()`.

From the performance point of view, it is worth knowing that most of the time in Lisp evaluation is spent executing `Lisp_Subr` and `Lisp_Compiled_Function` objects via `Ffuncall()` (not `Feval()`).

`Fapply()` implements Lisp `apply`, which is very similar to `funcall` except that if the last argument is a list, the result is the same as if each of the arguments in the list had been passed separately. `Fapply()` does some business to expand the last argument if it's a list, then calls `Ffuncall()` to do the work.

`apply1()`, `call0()`, `call1()`, `call2()`, and `call3()` call a function, passing it the argument(s) given (the arguments are given as separate C arguments rather than being passed as an array). `apply1()` uses `Fapply()` while the others use `Ffuncall()` to do the real work.

## 17.2 Dynamic Binding; The `specbinding` Stack; `Unwind-Protects`

```
struct specbinding
{
    Lisp_Object symbol;
    Lisp_Object old_value;
    Lisp_Object (*func) (Lisp_Object); /* for unwind-protect */
};
```

`struct specbinding` is used for local-variable bindings and `unwind-protects`. `specpdl` holds an array of `struct specbinding`'s, `specpdl_ptr` points to the beginning of the free bindings in the array, `specpdl_size` specifies the total number of binding slots in the array, and `max_specpdl_size` specifies the maximum number of bindings the array can be expanded to hold. `grow_specpdl()` increases the size of the `specpdl` array, multiplying its size by 2 but never exceeding `max_specpdl_size` (except that if this number is less than 400, it is first set to 400).

`specbind()` binds a symbol to a value and is used for local variables and `let` forms. The symbol and its old value (which might be `Unbound`, indicating no prior value) are recorded in the `specpdl` array, and `specpdl_size` is increased by 1.

`record_unwind_protect()` implements an `unwind-protect`, which, when placed around a section of code, ensures that some specified cleanup routine will be executed even if the code exits abnormally (e.g. through a `throw` or `quit`). `record_unwind_protect()` simply adds a new `specbinding` to the `specpdl` array and stores the appropriate information in it. The cleanup routine can either be a C function, which is stored in the `func` field, or a `progn` form, which is stored in the `old_value` field.

`unbind_to()` removes `specbindings` from the `specpdl` array until the specified position is reached. Each `specbinding` can be one of three types:

1. an `unwind-protect` with a C cleanup function (`func` is not 0, and `old_value` holds an argument to be passed to the function);
2. an `unwind-protect` with a Lisp form (`func` is 0, `symbol` is `nil`, and `old_value` holds the form to be executed with `Fprogn()`); or
3. a local-variable binding (`func` is 0, `symbol` is not `nil`, and `old_value` holds the old value, which is stored as the symbol's value).

## 17.3 Simple Special Forms

`or`, `and`, `if`, `cond`, `progn`, `prog1`, `prog2`, `setq`, `quote`, `function`, `let*`, `let`, `while`

All of these are very simple and work as expected, calling `Feval()` or `Fprogn()` as necessary and (in the case of `let` and `let*`) using `specbind()` to create bindings and `unbind_to()` to undo the bindings when finished.

Note that, with the exception of `Fprogn`, these functions are typically called in real life only in interpreted code, since the byte compiler knows how to convert calls to these functions directly into byte code.

## 17.4 Catch and Throw

```
struct catchtag
{
```

```

    Lisp_Object tag;
    Lisp_Object val;
    struct catchtag *next;
    struct gcpro *gcpro;
    jmp_buf jmp;
    struct backtrace *backlist;
    int lisp_eval_depth;
    int pdlcount;
};

```

`catch` is a Lisp function that places a catch around a body of code. A catch is a means of non-local exit from the code. When a catch is created, a tag is specified, and executing a `throw` to this tag will exit from the body of code caught with this tag, and its value will be the value given in the call to `throw`. If there is no such call, the code will be executed normally.

Information pertaining to a catch is held in a `struct catchtag`, which is placed at the head of a linked list pointed to by `catchlist`. `internal_catch()` is passed a C function to call (`Fprogn()` when Lisp `catch` is called) and arguments to give it, and places a catch around the function. Each `struct catchtag` is held in the stack frame of the `internal_catch()` instance that created the catch.

`internal_catch()` is fairly straightforward. It stores into the `struct catchtag` the tag name and the current values of `backtrace_list`, `lisp_eval_depth`, `gcprolist`, and the offset into the `specpdl` array, sets a jump point with `_setjmp()` (storing the jump point into the `struct catchtag`), and calls the function. Control will return to `internal_catch()` either when the function exits normally or through a `_longjmp()` to this jump point. In the latter case, `throw` will store the value to be returned into the `struct catchtag` before jumping. When it's done, `internal_catch()` removes the `struct catchtag` from the `catchlist` and returns the proper value.

`Fthrow()` goes up through the `catchlist` until it finds one with a matching tag. It then calls `unbind_catch()` to restore everything to what it was when the appropriate catch was set, stores the return value in the `struct catchtag`, and jumps (with `_longjmp()`) to its jump point.

`unbind_catch()` removes all catches from the `catchlist` until it finds the correct one. Some of the catches might have been placed for error-trapping, and if so, the appropriate entries on the `handlerlist` must be removed (see “errors”). `unbind_catch()` also restores the values of `gcprolist`, `backtrace_list`, and `lisp_eval`, and calls `unbind_to()` to undo any `specbindings` created since the catch.

## 18 Symbols and Variables

### 18.1 Introduction to Symbols

A symbol is basically just an object with four fields: a name (a string), a value (some Lisp object), a function (some Lisp object), and a property list (usually a list of alternating keyword/value pairs). What makes symbols special is that there is usually only one symbol with a given name, and the symbol is referred to by name. This makes a symbol a convenient way of calling up data by name, i.e. of implementing variables. (The variable's value is stored in the **value slot**.) Similarly, functions are referenced by name, and the definition of the function is stored in a symbol's **function slot**. This means that there can be a distinct function and variable with the same name. The property list is used as a more general mechanism of associating additional values with particular names, and once again the namespace is independent of the function and variable namespaces.

### 18.2 Obarrays

The identity of symbols with their names is accomplished through a structure called an obarray, which is just a poorly-implemented hash table mapping from strings to symbols whose name is that string. (I say “poorly implemented” because an obarray appears in Lisp as a vector with some hidden fields rather than as its own opaque type. This is an Emacs Lisp artifact that should be fixed.)

Obarrays are implemented as a vector of some fixed size (which should be a prime for best results), where each “bucket” of the vector contains one or more symbols, threaded through a hidden **next** field in the symbol. Lookup of a symbol in an obarray, and adding a symbol to an obarray, is accomplished through standard hash-table techniques.

The standard Lisp function for working with symbols and obarrays is **intern**. This looks up a symbol in an obarray given its name; if it's not found, a new symbol is automatically created with the specified name, added to the obarray, and returned. This is what happens when the Lisp reader encounters a symbol (or more precisely, encounters the name of a symbol) in some text that it is reading. There is a standard obarray called **obarray** that is used for this purpose, although the Lisp programmer is free to create his own obarrays and **intern** symbols in them.

Note that, once a symbol is in an obarray, it stays there until something is done about it, and the standard obarray **obarray** always stays around, so once you use any particular variable name, a corresponding symbol will stay around in **obarray** until you exit XEmacs.

Note that **obarray** itself is a variable, and as such there is a symbol in **obarray** whose name is "**obarray**" and which contains **obarray** as its value.

Note also that this call to **intern** occurs only when in the Lisp reader, not when the code is executed (at which point the symbol is already around, stored as such in the definition of the function).

You can create your own obarray using **make-vector** (this is horrible but is an artifact) and **intern** symbols into that obarray. Doing that will result in two or more symbols with the same name. However, at most one of these symbols is in the standard **obarray**: You cannot have two symbols of the same name in any particular obarray. Note that you cannot

add a symbol to an obarray in any fashion other than using `intern`: i.e. you can't take an existing symbol and put it in an existing obarray. Nor can you change the name of an existing symbol. (Since obarrays are vectors, you can violate the consistency of things by storing directly into the vector, but let's ignore that possibility.)

Usually symbols are created by `intern`, but if you really want, you can explicitly create a symbol using `make-symbol`, giving it some name. The resulting symbol is not in any obarray (i.e. it is `uninterned`), and you can't add it to any obarray. Therefore its primary purpose is as a symbol to use in macros to avoid namespace pollution. It can also be used as a carrier of information, but cons cells could probably be used just as well.

You can also use `intern-soft` to look up a symbol but not create a new one, and `unintern` to remove a symbol from an obarray. This returns the removed symbol. (Remember: You can't put the symbol back into any obarray.) Finally, `mapatoms` maps over all of the symbols in an obarray.

### 18.3 Symbol Values

The value field of a symbol normally contains a Lisp object. However, a symbol can be `unbound`, meaning that it logically has no value. This is internally indicated by storing a special Lisp object, called **the unbound marker** and stored in the global variable `Qunbound`. The unbound marker is of a special Lisp object type called `symbol-value-magic`. It is impossible for the Lisp programmer to directly create or access any object of this type.

You must not let any `\symbol-value-magic` object escape to the Lisp level. Printing any of these objects will cause the message `'INTERNAL EMACS BUG'` to appear as part of the print representation. (You may see this normally when you call `debug_print()` from the debugger on a Lisp object.) If you let one of these objects escape to the Lisp level, you will violate a number of assumptions contained in the C code and make the unbound marker not function right.

When a symbol is created, its value field (and function field) are set to `Qunbound`. The Lisp programmer can restore these conditions later using `makunbound` or `fmakunbound`, and can query to see whether the value of function fields are `bound` (i.e. have a value other than `Qunbound`) using `boundp` and `fboundp`. The fields are set to a normal Lisp object using `set` (or `setq`) and `fset`.

Other `symbol-value-magic` objects are used as special markers to indicate variables that have non-normal properties. This includes any variables that are tied into C variables (setting the variable magically sets some global variable in the C code, and likewise for retrieving the variable's value), variables that magically tie into slots in the current buffer, variables that are buffer-local, etc. The `symbol-value-magic` object is stored in the value cell in place of a normal object, and the code to retrieve a symbol's value (i.e. `symbol-value`) knows how to do special things with them. This means that you should not just fetch the value cell directly if you want a symbol's value.

The exact workings of this are rather complex and involved and are well-documented in comments in `'buffer.c'`, `'symbols.c'`, and `'lisp.h'`.

## 19 Buffers and Textual Representation

### 19.1 Introduction to Buffers

A buffer is logically just a Lisp object that holds some text. In this, it is like a string, but a buffer is optimized for frequent insertion and deletion, while a string is not. Furthermore:

1. Buffers are **permanent** objects, i.e. once you create them, they remain around, and need to be explicitly deleted before they go away.
2. Each buffer has a unique name, which is a string. Buffers are normally referred to by name. In this respect, they are like symbols.
3. Buffers have a default insertion position, called **point**. Inserting text (unless you explicitly give a position) goes at **point**, and moves **point** forward past the text. This is what is going on when you type text into Emacs.
4. Buffers have lots of extra properties associated with them.
5. Buffers can be **displayed**. What this means is that there exist a number of **windows**, which are objects that correspond to some visible section of your display, and each window has an associated buffer, and the current contents of the buffer are shown in that section of the display. The redisplay mechanism (which takes care of doing this) knows how to look at the text of a buffer and come up with some reasonable way of displaying this. Many of the properties of a buffer control how the buffer's text is displayed.
6. One buffer is distinguished and called the **current buffer**. It is stored in the variable `current_buffer`. Buffer operations operate on this buffer by default. When you are typing text into a buffer, the buffer you are typing into is always `current_buffer`. Switching to a different window changes the current buffer. Note that Lisp code can temporarily change the current buffer using `set-buffer` (often enclosed in a `save-excursion` so that the former current buffer gets restored when the code is finished). However, calling `set-buffer` will NOT cause a permanent change in the current buffer. The reason for this is that the top-level event loop sets `current_buffer` to the buffer of the selected window, each time it finishes executing a user command.

Make sure you understand the distinction between `current buffer` and `buffer of the selected window`, and the distinction between `point` of the current buffer and `window-point` of the selected window. (This latter distinction is explained in detail in the section on windows.)

### 19.2 The Text in a Buffer

The text in a buffer consists of a sequence of zero or more characters. A **character** is an integer that logically represents a letter, number, space, or other unit of text. Most of the characters that you will typically encounter belong to the ASCII set of characters, but there are also characters for various sorts of accented letters, special symbols, Chinese and Japanese ideograms (i.e. Kanji, Katakana, etc.), Cyrillic and Greek letters, etc. The actual number of possible characters is quite large.

For now, we can view a character as some non-negative integer that has some shape that defines how it typically appears (e.g. as an uppercase A). (The exact way in which a

character appears depends on the font used to display the character.) The internal type of characters in the C code is an `Ichar`; this is just an `int`, but using a symbolic type makes the code clearer.

Between every character in a buffer is a **buffer position** or **character position**. We can speak of the character before or after a particular buffer position, and when you insert a character at a particular position, all characters after that position end up at new positions. When we speak of the character **at** a position, we really mean the character after the position. (This schizophrenia between a buffer position being “between” two characters and “on” a character is rampant in Emacs.)

Buffer positions are numbered starting at 1. This means that position 1 is before the first character, and position 0 is not valid. If there are  $N$  characters in a buffer, then buffer position  $N+1$  is after the last one, and position  $N+2$  is not valid.

The internal makeup of the `Ichar` integer varies depending on whether we have compiled with MULE support. If not, the `Ichar` integer is an 8-bit integer with possible values from 0 - 255. 0 - 127 are the standard ASCII characters, while 128 - 255 are the characters from the ISO-8859-1 character set. If we have compiled with MULE support, an `Ichar` is a 19-bit integer, with the various bits having meanings according to a complex scheme that will be detailed later. The characters numbered 0 - 255 still have the same meanings as for the non-MULE case, though.

Internally, the text in a buffer is represented in a fairly simple fashion: as a contiguous array of bytes, with a **gap** of some size in the middle. Although the gap is of some substantial size in bytes, there is no text contained within it: From the perspective of the text in the buffer, it does not exist. The gap logically sits at some buffer position, between two characters (or possibly at the beginning or end of the buffer). Insertion of text in a buffer at a particular position is always accomplished by first moving the gap to that position (i.e. through some block moving of text), then writing the text into the beginning of the gap, thereby shrinking the gap. If the gap shrinks down to nothing, a new gap is created. (What actually happens is that a new gap is “created” at the end of the buffer’s text, which requires nothing more than changing a couple of indices; then the gap is “moved” to the position where the insertion needs to take place by moving up in memory all the text after that position.) Similarly, deletion occurs by moving the gap to the place where the text is to be deleted, and then simply expanding the gap to include the deleted text. (**Expanding** and **shrinking** the gap as just described means just that the internal indices that keep track of where the gap is located are changed.)

Note that the total amount of memory allocated for a buffer text never decreases while the buffer is live. Therefore, if you load up a 20-megabyte file and then delete all but one character, there will be a 20-megabyte gap, which won’t get any smaller (except by inserting characters back again). Once the buffer is killed, the memory allocated for the buffer text will be freed, but it will still be sitting on the heap, taking up virtual memory, and will not be released back to the operating system. (However, if you have compiled XEmacs with `rel-alloc`, the situation is different. In this case, the space *will* be released back to the operating system. However, this tends to result in a noticeable speed penalty.)

Astute readers may notice that the text in a buffer is represented as an array of *bytes*, while (at least in the MULE case) an `Ichar` is a 19-bit integer, which clearly cannot fit in a byte. This means (of course) that the text in a buffer uses a different representation from

an Ichar: specifically, the 19-bit Ichar becomes a series of one to four bytes. The conversion between these two representations is complex and will be described later.

In the non-MULE case, everything is very simple: An Ichar is an 8-bit value, which fits neatly into one byte.

If we are given a buffer position and want to retrieve the character at that position, we need to follow these steps:

1. Pretend there's no gap, and convert the buffer position into a **byte index** that indexes to the appropriate byte in the buffer's stream of textual bytes. By convention, byte indices begin at 1, just like buffer positions. In the non-MULE case, byte indices and buffer positions are identical, since one character equals one byte.
2. Convert the byte index into a **memory index**, which takes the gap into account. The memory index is a direct index into the block of memory that stores the text of a buffer. This basically just involves checking to see if the byte index is past the gap, and if so, adding the size of the gap to it. By convention, memory indices begin at 1, just like buffer positions and byte indices, and when referring to the position that is **at** the gap, we always use the memory position at the *beginning*, not at the end, of the gap.
3. Fetch the appropriate bytes at the determined memory position.
4. Convert these bytes into an Ichar.

In the non-Mule case, (3) and (4) boil down to a simple one-byte memory access.

Note that we have defined three types of positions in a buffer:

1. **bu er positions** or **character positions**, typedef Charbpos
2. **byte indices**, typedef Bytebpos
3. **memory indices**, typedef Membpos

All three typedefs are just **ints**, but defining them this way makes things a lot clearer.

Most code works with buffer positions. In particular, all Lisp code that refers to text in a buffer uses buffer positions. Lisp code does not know that byte indices or memory indices exist.

Finally, we have a typedef for the bytes in a buffer. This is a **Ibyte**, which is an unsigned char. Referring to them as Ibytes underscores the fact that we are working with a string of bytes in the internal Emacs buffer representation rather than in one of a number of possible alternative representations (e.g. EUC-encoded text, etc.).

### 19.3 Buffer Lists

Recall earlier that buffers are **permanent** objects, i.e. that they remain around until explicitly deleted. This entails that there is a list of all the buffers in existence. This list is actually an assoc-list (mapping from the buffer's name to the buffer) and is stored in the global variable `Vbuffer_alist`.

The order of the buffers in the list is important: the buffers are ordered approximately from most-recently-used to least-recently-used. Switching to a buffer using `switch-to-buffer`, `pop-to-buffer`, etc. and switching windows using `other-window`, etc. usually brings the new current buffer to the front of the list. `switch-to-buffer`, `other-buffer`, etc. look at the beginning of the list to find an alternative buffer to suggest. You can also explicitly move a buffer to the end of the list using `bury-buffer`.

In addition to the global ordering in `Vbuffer_alist`, each frame has its own ordering of the list. These lists always contain the same elements as in `Vbuffer_alist` although possibly in a different order. `buffer-list` normally returns the list for the selected frame. This allows you to work in separate frames without things interfering with each other.

The standard way to look up a buffer given a name is `get-buffer`, and the standard way to create a new buffer is `get-buffer-create`, which looks up a buffer with a given name, creating a new one if necessary. These operations correspond exactly with the symbol operations `intern-soft` and `intern`, respectively. You can also force a new buffer to be created using `generate-new-buffer`, which takes a name and (if necessary) makes a unique name from this by appending a number, and then creates the buffer. This is basically like the symbol operation `gensym`.

## 19.4 Markers and Extents

Among the things associated with a buffer are things that are logically attached to certain buffer positions. This can be used to keep track of a buffer position when text is inserted and deleted, so that it remains at the same spot relative to the text around it; to assign properties to particular sections of text; etc. There are two such objects that are useful in this regard: they are **markers** and **extents**.

A **marker** is simply a flag placed at a particular buffer position, which is moved around as text is inserted and deleted. Markers are used for all sorts of purposes, such as the `mark` that is the other end of textual regions to be cut, copied, etc.

An **extent** is similar to two markers plus some associated properties, and is used to keep track of regions in a buffer as text is inserted and deleted, and to add properties (e.g. fonts) to particular regions of text. The external interface of extents is explained elsewhere.

The important thing here is that markers and extents simply contain buffer positions in them as integers, and every time text is inserted or deleted, these positions must be updated. In order to minimize the amount of shuffling that needs to be done, the positions in markers and extents (there's one per marker, two per extent) are stored in `Mempos`'s. This means that they only need to be moved when the text is physically moved in memory; since the gap structure tries to minimize this, it also minimizes the number of marker and extent indices that need to be adjusted. Look in `'insdel.c'` for the details of how this works.

One other important distinction is that markers are **temporary** while extents are **permanent**. This means that markers disappear as soon as there are no more pointers to them, and correspondingly, there is no way to determine what markers are in a buffer if you are just given the buffer. Extents remain in a buffer until they are detached (which could happen as a result of text being deleted) or the buffer is deleted, and primitives do exist to enumerate the extents in a buffer.

## 19.5 Ibytes and Ichars

Not yet documented.

## 19.6 The Buffer Object

Buffers contain fields not directly accessible by the Lisp programmer. We describe them here, naming them by the names used in the C code. Many are accessible indirectly in Lisp programs via Lisp primitives.

- name**           The buffer name is a string that names the buffer. It is guaranteed to be unique. See section “Buffer Names” in [XEmacs Lisp Reference Manual](#)
- save\_modified**  
This field contains the time when the buffer was last saved, as an integer. See section “Buffer Modification” in [XEmacs Lisp Reference Manual](#)
- modtime**       This field contains the modification time of the visited file. It is set when the file is written or read. Every time the buffer is written to the file, this field is compared to the modification time of the file. See section “Buffer Modification” in [XEmacs Lisp Reference Manual](#)
- auto\_save\_modified**  
This field contains the time when the buffer was last auto-saved.
- last\_window\_start**  
This field contains the `window-start` position in the buffer as of the last time the buffer was displayed in a window.
- undo\_list**  
This field points to the buffer’s undo list. See section “Undo” in [XEmacs Lisp Reference Manual](#)
- syntax\_table\_v**  
This field contains the syntax table for the buffer. See section “Syntax Tables” in [XEmacs Lisp Reference Manual](#)
- downcase\_table**  
This field contains the conversion table for converting text to lower case. See section “Case Tables” in [XEmacs Lisp Reference Manual](#)
- upcase\_table**  
This field contains the conversion table for converting text to upper case. See section “Case Tables” in [XEmacs Lisp Reference Manual](#)
- case\_canon\_table**  
This field contains the conversion table for canonicalizing text for case-folding search. See section “Case Tables” in [XEmacs Lisp Reference Manual](#)
- case\_eqv\_table**  
This field contains the equivalence table for case-folding search. See section “Case Tables” in [XEmacs Lisp Reference Manual](#)
- display\_table**  
This field contains the buffer’s display table, or `nil` if it doesn’t have one. See section “Display Tables” in [XEmacs Lisp Reference Manual](#)
- markers**       This field contains the chain of all markers that currently point into the buffer. Deletion of text in the buffer, and motion of the buffer’s gap, must check each

of these markers and perhaps update it. See section “Markers” in [XEmacs Lisp Reference Manual](#)

**backed\_up**

This field is a flag that tells whether a backup file has been made for the visited file of this buffer.

**mark**

This field contains the mark for the buffer. The mark is a marker, hence it is also included on the list `markers`. See section “The Mark” in [XEmacs Lisp Reference Manual](#)

**mark\_active**

This field is non-`nil` if the buffer’s mark is active.

**local\_var\_alist**

This field contains the association list describing the variables local in this buffer, and their values, with the exception of local variables that have special slots in the buffer object. (Those slots are omitted from this table.) See section “Buffer-Local Variables” in [XEmacs Lisp Reference Manual](#)

**modeline\_format**

This field contains a Lisp object which controls how to display the mode line for this buffer. See section “Modeline Format” in [XEmacs Lisp Reference Manual](#)

**base\_buffer**

This field holds the buffer’s base buffer (if it is an indirect buffer), or `nil`.

## 19.7 Searching and Matching

Very incomplete, limited to a brief introduction.

People find the searching and matching code difficult to understand. And indeed, the details are hard. However, the basic structures are not so complex. First, there’s a hard question with a simple answer. What about Mule? The answer here is that it turns out that Mule characters can be matched byte by byte, so neither the search code nor the regular expression code need take much notice of it at all! Of course, we add some special features (such as regular expressions that match only certain charsets), but these do not require new concepts. The main exception is that wild-card matches in Mule have to be careful to swallow whole characters. This is handled using the same basic macros that are used for buffer and string movements.

The complex algorithms for searching are for simple string searches. In particular, the algorithm used for fast string searching is Boyer-Moore. This algorithm is based on the idea that if you have a mismatch at a given position, you can precompute where to restart the search. This typically means that you can often make many fewer than  $N$  character comparisons, where  $N$  is the position at which the match is found, or the size of the text if it contains no match. That’s fast! But it’s not easy. You must “compile” the search string into a jump table. See the source, ‘`search.c`’, for more information.

Emacs changes the basic algorithms somewhat in order to handle case-insensitive searches without a full-blown regular expression.

Regular expressions, on the other hand, have a trivial search implementation: try a match at each position. (Under POSIX rules, it’s a bit more complex, because POSIX

requires that you find the *longest* match in the text. This means you keep a record of the best match so far, and find all the matches.)

The matching code for regular expressions is quite complex. First, the regular expression itself is compiled. There are two basic approaches that could be taken. The first is to compile the expression into tables to drive a generic finite automaton emulator. This is the approach given in many textbooks (Sedgewick's *Algorithms* and Aho, Sethi, and Ullmann's *Compilers: Principles, Techniques, and Tools*, aka "The Dragon Book") as well as being used by the 'lex' family of lexical analysis engines.

Emacs uses a somewhat different technique. The expression is compiled into a form of bytecode, which is interpreted by a special interpreter. The interpreter itself basically amounts to an inline implementation of the finite automaton emulator. The advantage of this technique is that it's easier to add special features, such as control of case-sensitivity via a global variable.

The compiler is not treated here. See the source, 'regex.c'. The interpreter, although it is divided into several functions, and looks fearsomely complex, is actually quite simple in concept. However, basically what you're doing there is a strcmp on steroids, right?

```
int
strcmp (char *p,          /* pattern pointer */
        char *b)         /* buffer pointer */
{
    while (*p++ == *b++)
        ;
    return * (--p) - * (--b); /* oops, we overshoot */
}
```

Really, it's no harder than that. (A bit of a white lie, OK?)

How does the regexp code generalize this?

1. Depending on the pattern, **\*b** may have a general relationship to **\*p**. *I.e.*, direct comparison against **\*p** is generalized to include checks for set membership, and context dependent properties. This depends on **&\*b**. Of course that's meaningless in C, so we use **b** directly, instead.
2. Although to ensure the algorithm terminates, **b** must advance step by step, **p** can branch and jump.
3. The information returned is much greater, including information about subexpressions.

We'll ignore (3). (2) is mostly interesting when compiling the regular expression. Now we have

```
enum operator_t {
    accept = 0,
    exact,
    any,
    range,
    group,      /* actually, these are probably */
    repeat,     /* turned into conditional code */
    /* etc */
};
```

```

enum status_t {
    working = 0,
    matched,
    mismatch,
    end_of_buffer,
    error
};

struct pattern {
    enum operator_t operator;
    char char_value;
    boolean range_table[256];
    /* etc, etc */
};

char *p, /* pattern pointer */
      *b; /* buffer pointer */

enum status_t
match (struct pattern *p, char *b)
{
    enum status_t done = working;

    while (!(done = match_1_operator (p, b)))
    {
        struct pattern *p1 = p;
        p = next_p (p, b);
        b = next_b (p1, b);
    }
    return done;
}

```

This format exposes the underlying finite automaton.

All of them have the following structure, except that the ‘next\_’ functions decide where to jump (for ‘p’) and whether or not to increment (for ‘b’), rather than checking for satisfaction of a matching condition.

```

enum status_t
match_1_operator (pattern *p, char *b)
{
    if (! *b) return end_of_buffer;
    switch (p->operator)
    {
        case accept:
            return matched;
        case exact:
            if (*b != p->char_value) return mismatch; else break;
    }
}

```

```

    case any:
        break;
    case range:
        /* range_table is computed in the regexp_compile function */
        if (! p->range_table[*b]) return mismatch;
        /* etc, etc */
    }
    return working;
}

```

Grouping, repetition, and alternation are handled by compiling the subexpression and calling `match (p->subpattern, b)` recursively.

In terms of reading the actual code, there are five optimizations (obfuscations, if you like) that have been done.

1. An explicit "failure stack" has been substituted for recursion.
2. The `match_1_operator`, `next_p`, and `next_b` functions are actually inlined into the `match` function for efficiency. Then the pointer movement is interspersed with the matching operations.
3. If the operator uses buffer context, the buffer pointer movement is sometimes implicit in the operations retrieving the context.
4. Some cases are combined into short preparation for individual cases, and a "fall-through" into combined code for several cases.
5. The `pattern` type is not an explicit `'struct'`. Instead, the data (including, *e.g.*, `'range_table'`) is inlined into the compiled bytecode. This leads to bizarre code in the interpreter like

```

    case range:
        p += *(p + 1); break;
in next_p, because the compiled pattern is laid out
    ..., 'range', count, first_8_flags, second_8_flags, ..., next_op, ...

```

But if you keep your eye on the "switch in a loop" structure, you should be able to understand the parts you need.



## 20 MULE Character Sets and Encodings

Recall that there are two primary ways that text is represented in XEmacs. The `buffer` representation sees the text as a series of bytes (Ibytes), with a variable number of bytes used per character. The `character` representation sees the text as a series of integers (Ichars), one per character. The character representation is a cleaner representation from a theoretical standpoint, and is thus used in many cases when lots of manipulations on a string need to be done. However, the buffer representation is the standard representation used in both Lisp strings and buffers, and because of this, it is the “default” representation that text comes in. The reason for using this representation is that it’s compact and is compatible with ASCII.

### 20.1 Character Sets

A character set (or `charset`) is an ordered set of characters. A particular character in a charset is indexed using one or more `position codes` which are non-negative integers. The number of position codes needed to identify a particular character in a charset is called the `dimension` of the charset. In XEmacs/Mule, all charsets have dimension 1 or 2, and the size of all charsets (except for a few special cases) is either 94, 96, 94 by 94, or 96 by 96. The range of position codes used to index characters from any of these types of character sets is as follows:

Charset type	Position code 1	Position code 2
94	33 - 126	N/A
96	32 - 127	N/A
94x94	33 - 126	33 - 126
96x96	32 - 127	32 - 127

Note that in the above cases position codes do not start at an expected value such as 0 or 1. The reason for this will become clear later.

For example, Latin-1 is a 96-character charset, and JISX0208 (the Japanese national character set) is a 94x94-character charset.

[Note that, although the ranges above define the *valid* position codes for a charset, some of the slots in a particular charset may in fact be empty. This is the case for JISX0208, for example, where (e.g.) all the slots whose first position code is in the range 118 - 127 are empty.]

There are three charsets that do not follow the above rules. All of them have one dimension, and have ranges of position codes as follows:

Charset name	Position code 1
ASCII	0 - 127
Control-1	0 - 31
Composite	0 - some large number

(The upper bound of the position code for composite characters has not yet been determined, but it will probably be at least 16,383).

ASCII is the union of two subsidiary character sets: Printing-ASCII (the printing ASCII character set, consisting of position codes 33 - 126, like for a standard 94-character charset)

and Control-ASCII (the non-printing characters that would appear in a binary file with codes 0 - 32 and 127).

Control-1 contains the non-printing characters that would appear in a binary file with codes 128 - 159.

Composite contains characters that are generated by overstriking one or more characters from other charsets.

Note that some characters in ASCII, and all characters in Control-1, are **control** (non-printing) characters. These have no printed representation but instead control some other function of the printing (e.g. TAB or 8 moves the current character position to the next tab stop). All other characters in all charsets are **graphic** (printing) characters.

When a binary file is read in, the bytes in the file are assigned to character sets as follows:

Bytes	Character set	Range
0 - 127	ASCII	0 - 127
128 - 159	Control-1	0 - 31
160 - 255	Latin-1	32 - 127

This is a bit ad-hoc but gets the job done.

## 20.2 Encodings

An **encoding** is a way of numerically representing characters from one or more character sets. If an encoding only encompasses one character set, then the position codes for the characters in that character set could be used directly. This is not possible, however, if more than one character set is to be used in the encoding.

For example, the conversion detailed above between bytes in a binary file and characters is effectively an encoding that encompasses the three character sets ASCII, Control-1, and Latin-1 in a stream of 8-bit bytes.

Thus, an encoding can be viewed as a way of encoding characters from a specified group of character sets using a stream of bytes, each of which contains a fixed number of bits (but not necessarily 8, as in the common usage of “byte”).

Here are descriptions of a couple of common encodings:

### 20.2.1 Japanese EUC (Extended Unix Code)

This encompasses the character sets Printing-ASCII, Japanese-JISX0201, and Japanese-JISX0208-Kana (half-width katakana, the right half of JISX0201). It uses 8-bit bytes.

Note that Printing-ASCII and Japanese-JISX0201-Kana are 94-character charsets, while Japanese-JISX0208 is a 94x94-character charset.

The encoding is as follows:

Character set	Representation (PC=position-code)
Printing-ASCII	PC1
Japanese-JISX0201-Kana	0x8E   PC1 + 0x80
Japanese-JISX0208	PC1 + 0x80   PC2 + 0x80
Japanese-JISX0212	PC1 + 0x80   PC2 + 0x80

## 20.2.2 JIS7

This encompasses the character sets Printing-ASCII, Japanese-JISX0201-Roman (the left half of JISX0201; this character set is very similar to Printing-ASCII and is a 94-character charset), Japanese-JISX0208, and Japanese-JISX0201-Kana. It uses 7-bit bytes.

Unlike Japanese EUC, this is a **modal** encoding, which means that there are multiple states that the encoding can be in, which affect how the bytes are to be interpreted. Special sequences of bytes (called **escape sequences**) are used to change states.

The encoding is as follows:

Character set	Representation (PC=position-code)
-----	-----
Printing-ASCII	PC1
Japanese-JISX0201-Roman	PC1
Japanese-JISX0201-Kana	PC1
Japanese-JISX0208	PC1 PC2

Escape sequence	ASCII equivalent	Meaning
-----	-----	-----
0x1B 0x28 0x4A	ESC ( J	invoke Japanese-JISX0201-Roman
0x1B 0x28 0x49	ESC ( I	invoke Japanese-JISX0201-Kana
0x1B 0x24 0x42	ESC \$ B	invoke Japanese-JISX0208
0x1B 0x28 0x42	ESC ( B	invoke Printing-ASCII

Initially, Printing-ASCII is invoked.

## 20.3 Internal Mule Encodings

In XEmacs/Mule, each character set is assigned a unique number, called a **leading byte**. This is used in the encodings of a character. Leading bytes are in the range 0x80 - 0xFF (except for ASCII, which has a leading byte of 0), although some leading bytes are reserved.

Charsets whose leading byte is in the range 0x80 - 0x9F are called **official** and are used for built-in charsets. Other charsets are called **private** and have leading bytes in the range 0xA0 - 0xFF; these are user-defined charsets.

More specifically:

Character set	Leading byte
-----	-----
ASCII	0
Composite	0x80
Dimension-1 Official	0x81 - 0x8D (0x8E is free)
Control-1	0x8F
Dimension-2 Official	0x90 - 0x99 (0x9A - 0x9D are free; 0x9E and 0x9F are reserved)
Dimension-1 Private	0xA0 - 0xEF
Dimension-2 Private	0xF0 - 0xFF

There are two internal encodings for characters in XEmacs/Mule. One is called **string encoding** and is an 8-bit encoding that is used for representing characters in a buffer or string. It uses 1 to 4 bytes per character. The other is called **character encoding** and is a 19-bit encoding that is used for representing characters individually in a variable.

(In the following descriptions, we'll ignore composite characters for the moment. We also give a general (structural) overview first, followed later by the exact details.)

### 20.3.1 Internal String Encoding

ASCII characters are encoded using their position code directly. Other characters are encoded using their leading byte followed by their position code(s) with the high bit set. Characters in private character sets have their leading byte prefixed with a **leading byte pre x**, which is either 0x9E or 0x9F. (No character sets are ever assigned these leading bytes.) Specifically:

Character set -----	Encoding (PC=position-code, LB=leading-byte) -----
ASCII	PC-1
Control-1	LB   PC1 + 0xA0
Dimension-1 official	LB   PC1 + 0x80
Dimension-1 private	0x9E   LB   PC1 + 0x80
Dimension-2 official	LB   PC1 + 0x80   PC2 + 0x80
Dimension-2 private	0x9F   LB   PC1 + 0x80   PC2 + 0x80

The basic characteristic of this encoding is that the first byte of all characters is in the range 0x00 - 0x9F, and the second and following bytes of all characters is in the range 0xA0 - 0xFF. This means that it is impossible to get out of sync, or more specifically:

1. Given any byte position, the beginning of the character it is within can be determined in constant time.
2. Given any byte position at the beginning of a character, the beginning of the next character can be determined in constant time.
3. Given any byte position at the beginning of a character, the beginning of the previous character can be determined in constant time.
4. Textual searches can simply treat encoded strings as if they were encoded in a one-byte-per-character fashion rather than the actual multi-byte encoding.

None of the standard non-modal encodings meet all of these conditions. For example, EUC satisfies only (2) and (3), while Shift-JIS and Big5 (not yet described) satisfy only (2). (All non-modal encodings must satisfy (2), in order to be unambiguous.)

### 20.3.2 Internal Character Encoding

One 19-bit word represents a single character. The word is separated into three fields:

Bit number:	18 17 16 15 14 13 12 11 10 09 08 07 06 05 04 03 02 01 00
	<-----> <-----> <----->
Field:	1 2 3

Note that fields 2 and 3 hold 7 bits each, while field 1 holds 5 bits.

Character set -----	Field 1 -----	Field 2 -----	Field 3 -----
------------------------	------------------	------------------	------------------

ASCII	0	0	PC1
range:			(00 - 7F)
Control-1	0	1	PC1
range:			(00 - 1F)
Dimension-1 official	0	LB - 0x80	PC1
range:		(01 - 0D)	(20 - 7F)
Dimension-1 private	0	LB - 0x80	PC1
range:		(20 - 6F)	(20 - 7F)
Dimension-2 official	LB - 0x8F	PC1	PC2
range:	(01 - 0A)	(20 - 7F)	(20 - 7F)
Dimension-2 private	LB - 0xE1	PC1	PC2
range:	(0F - 1E)	(20 - 7F)	(20 - 7F)
Composite	0x1F	?	?

Note that character codes 0 - 255 are the same as the “binary encoding” described above.

## 20.4 CCL

CCL PROGRAM SYNTAX:

```
CCL_PROGRAM := (CCL_MAIN_BLOCK
                [ CCL_EOF_BLOCK ])
```

```
CCL_MAIN_BLOCK := CCL_BLOCK
```

```
CCL_EOF_BLOCK := CCL_BLOCK
```

```
CCL_BLOCK := STATEMENT | (STATEMENT [STATEMENT ...])
```

```
STATEMENT :=
```

```
  SET | IF | BRANCH | LOOP | REPEAT | BREAK
  | READ | WRITE
```

```
SET := (REG = EXPRESSION) | (REG SELF_OP EXPRESSION)
      | INT-OR-CHAR
```

```
EXPRESSION := ARG | (EXPRESSION OP ARG)
```

```
IF := (if EXPRESSION CCL_BLOCK CCL_BLOCK)
```

```
BRANCH := (branch EXPRESSION CCL_BLOCK [CCL_BLOCK ...])
```

```
LOOP := (loop STATEMENT [STATEMENT ...])
```

```
BREAK := (break)
```

```
REPEAT := (repeat)
```

```
  | (write-repeat [REG | INT-OR-CHAR | string])
```

```
  | (write-read-repeat REG [INT-OR-CHAR | string | ARRAY]?)
```

```
READ := (read REG) | (read REG REG)
```

```
  | (read-if REG ARITH_OP ARG CCL_BLOCK CCL_BLOCK)
```

```
  | (read-branch REG CCL_BLOCK [CCL_BLOCK ...])
```

```
WRITE := (write REG) | (write REG REG)
```

```
  | (write INT-OR-CHAR) | (write STRING) | STRING
```

```
  | (write REG ARRAY)
```

```

END := (end)

REG := r0 | r1 | r2 | r3 | r4 | r5 | r6 | r7
ARG := REG | INT-OR-CHAR
OP :=  + | - | * | / | % | & | '|' | ^ | << | >> | <8 | >8 | //
      | < | > | == | <= | >= | !=
SELF_OP :=
      += | -= | *= | /= | %= | &= | |= | ^= | <<= | >>=
ARRAY := '[' INT-OR-CHAR ... ']'
INT-OR-CHAR := INT | CHAR

```

#### MACHINE CODE:

The machine code consists of a vector of 32-bit words. The first such word specifies the start of the EOF section of the code; this is the code executed to handle any stuff that needs to be done (e.g. designating back to ASCII and left-to-right mode) after all other encoded/decoded data has been written out. This is not used for charset CCL programs.

REGISTER: 0..7 -- referred by RRR or rrr

OPERATOR BIT FIELD (27-bit): XXXXXXXXXXXXXXXX RRR TTTT

TTTT (5-bit): operator type

RRR (3-bit): register number

XXXXXXXXXXXXXXXXXX (15-bit):

CCCCCCCCCCCCCCC: constant or address

00000000000rrr: register number

```

AAAA:  00000 +
        00001 -
        00010 *
        00011 /
        00100 %
        00101 &
        00110 |
        00111 ~

        01000 <<
        01001 >>
        01010 <8
        01011 >8
        01100 //
        01101 not used
        01110 not used
        01111 not used

```

	10000 <	
	10001 >	
	10010 ==	
	10011 <=	
	10100 >=	
	10101 !=	
OPERATORS:	TTTTT RRR XX..	
SetCS:	00000 RRR C...C	RRR = C...C
SetCL:	00001 RRR ..... c.....C	RRR = c...c
SetR:	00010 RRR ..rrr	RRR = rrr
SetA:	00011 RRR ..rrr C.....C c.....c	RRR = array[rrr] size of array = C...C contents = c...c
Jump:	00100 000 c...c	jump to c...c
JumpCond:	00101 RRR c...c	if (!RRR) jump to c...c
WriteJump:	00110 RRR c...c	Write1 RRR, jump to c...c
WriteReadJump:	00111 RRR c...c	Write1, Read1 RRR, jump to c...c
WriteCJump:	01000 000 c...c C...C	Write1 C...C, jump to c...c
WriteCReadJump:	01001 RRR c...c C.....C	Write1 C...C, Read1 RRR, and jump to c...c
WriteSJump:	01010 000 c...c C.....C S.....S ...	WriteS, jump to c...c
WriteSReadJump:	01011 RRR c...c C.....C S.....S ...	WriteS, Read1 RRR, jump to c...c
WriteAReadJump:	01100 RRR c...c C.....C c.....c ...	WriteA, Read1 RRR, jump to c...c size of array = C...C contents = c...c
Branch:	01101 RRR C...C c.....C	if (RRR >= 0 && RRR < C..) branch to (RRR+1)th address
Read1:	01110 RRR ...	read 1-byte to RRR
Read2:	01111 RRR ..rrr	read 2-byte to RRR and rrr
ReadBranch:	10000 RRR C...C c.....C ...	Read1 and Branch
Write1:	10001 RRR ..... Write2:	write 1-byte RRR
Write2:	10010 RRR ..rrr	write 2-byte RRR and rrr
WriteC:	10011 000 ..... WriteC:	write 1-char C...CC

	C.....C	
WriteS:	10100 000 .....	write C..-byte of string
	C.....C	
	S.....S	
	...	
WriteA:	10101 RRR .....	write array[RRR]
	C.....C	size of array = C...C
	c.....c	contents = c...c
	...	
End:	10110 000 .....	terminate the execution
SetSelfCS:	10111 RRR C...C	RRR AAAAA= C...C
	.....AAAAA	
SetSelfCL:	11000 RRR .....	RRR AAAAA= c...c
	c.....c	
	.....AAAAA	
SetSelfR:	11001 RRR ..Rrr	RRR AAAAA= rrr
	.....AAAAA	
SetExprCL:	11010 RRR ..Rrr	RRR = rrr AAAAA c...c
	c.....c	
	.....AAAAA	
SetExprR:	11011 RRR ..rrr	RRR = rrr AAAAA Rrr
	.....Rrr	
	.....AAAAA	
JumpCondC:	11100 RRR c...c	if !(RRR AAAAA C..) jump to c...c
	C.....C	
	.....AAAAA	
JumpCondR:	11101 RRR c...c	if !(RRR AAAAA rrr) jump to c...c
	.....rrr	
	.....AAAAA	
ReadJumpCondC:	11110 RRR c...c	Read1 and JumpCondC
	C.....C	
	.....AAAAA	
ReadJumpCondR:	11111 RRR c...c	Read1 and JumpCondR
	.....rrr	
	.....AAAAA	

## 21 The Lisp Reader and Compiler

Not yet documented.



## 22 Lstreams

An **lstream** is an internal Lisp object that provides a generic buffering stream implementation. Conceptually, you send data to the stream or read data from the stream, not caring what's on the other end of the stream. The other end could be another stream, a file descriptor, a stdio stream, a fixed block of memory, a reallocating block of memory, etc. The main purpose of the stream is to provide a standard interface and to do buffering. Macros are defined to read or write characters, so the calling functions do not have to worry about blocking data together in order to achieve efficiency.

### 22.1 Creating an Lstream

Lstreams come in different types, depending on what is being interfaced to. Although the primitive for creating new lstreams is `Lstream_new()`, generally you do not call this directly. Instead, you call some type-specific creation function, which creates the lstream and initializes it as appropriate for the particular type.

All lstream creation functions take a **mode** argument, specifying what mode the lstream should be opened as. This controls whether the lstream is for input and output, and optionally whether data should be blocked up in units of MULE characters. Note that some types of lstreams can only be opened for input; others only for output; and others can be opened either way. #### Richard Mlynarik thinks that there should be a strict separation between input and output streams, and he's probably right.

**mode** is a string, one of

"r"	Open for reading.
"w"	Open for writing.
"rc"	Open for reading, but "read" never returns partial MULE characters.
"wc"	Open for writing, but never writes partial MULE characters.

### 22.2 Lstream Types

stdio

filedesc

lisp-string

fixed-buffer

resizing-buffer

dynarr

lisp-buffer

print

decoding

encoding

## 22.3 Lstream Functions

**Lstream \*** `Lstream_new` (`Lstream_implementation *imp`, `const char *mode`) [Function]

Allocate and return a new Lstream. This function is not really meant to be called directly; rather, each stream type should provide its own stream creation function, which creates the stream and does any other necessary creation stuff (e.g. opening a file).

**void** `Lstream_set_buffering` (`Lstream *lstr`, `Lstream_buffering buffering`, `int buffering_size`) [Function]

Change the buffering of a stream. See ‘`lstream.h`’. By default the buffering is `STREAM_BLOCK_BUFFERED`.

**int** `Lstream_flush` (`Lstream *lstr`) [Function]

Flush out any pending unwritten data in the stream. Clear any buffered input data. Returns 0 on success, -1 on error.

**int** `Lstream_putc` (`Lstream *stream`, `int c`) [Macro]

Write out one byte to the stream. This is a macro and so it is very efficient. The `c` argument is only evaluated once but the `stream` argument is evaluated more than once. Returns 0 on success, -1 on error.

**int** `Lstream_getc` (`Lstream *stream`) [Macro]

Read one byte from the stream. This is a macro and so it is very efficient. The `stream` argument is evaluated more than once. Return value is -1 for EOF or error.

**void** `Lstream_ungetc` (`Lstream *stream`, `int c`) [Macro]

Push one byte back onto the input queue. This will be the next byte read from the stream. Any number of bytes can be pushed back and will be read in the reverse order they were pushed back—most recent first. (This is necessary for consistency—if there are a number of bytes that have been unread and I read and unread a byte, it needs to be the first to be read again.) This is a macro and so it is very efficient. The `c` argument is only evaluated once but the `stream` argument is evaluated more than once.

**int** `Lstream_fputc` (`Lstream *stream`, `int c`) [Function]

**int** `Lstream_fgetc` (`Lstream *stream`) [Function]

**void** `Lstream_fungetc` (`Lstream *stream`, `int c`) [Function]

Function equivalents of the above macros.

**Bytecount** `Lstream_read` (`Lstream *stream`, `void *data`, `Bytecount size`) [Function]

Read `size` bytes of `data` from the stream. Return the number of bytes read. 0 means EOF. -1 means an error occurred and no bytes were read.

**Bytecount** `Lstream_write` (`Lstream *stream`, `void *data`, `Bytecount size`) [Function]

Write `size` bytes of `data` to the stream. Return the number of bytes written. -1 means an error occurred and no bytes were written.

**void Lstream\_unread** (Lstream \*stream, void \*data, Bytecount size) [Function]  
 Push back **size** bytes of **data** onto the input queue. The next call to **Lstream\_read()** with the same **size** will read the same bytes back. Note that this will be the case even if there is other pending unread data.

**int Lstream\_close** (Lstream \*stream) [Function]  
 Close the stream. All data will be flushed out.

**void Lstream\_reopen** (Lstream \*stream) [Function]  
 Reopen a closed stream. This enables I/O on it again. This is not meant to be called except from a wrapper routine that reinitializes variables and such—the close routine may well have freed some necessary storage structures, for example.

**void Lstream\_rewind** (Lstream \*stream) [Function]  
 Rewind the stream to the beginning.

## 22.4 Lstream Methods

**Bytecount reader** (Lstream \*stream, unsigned char \*data, Bytecount size) [Lstream Method]

Read some data from the stream's end and store it into **data**, which can hold **size** bytes. Return the number of bytes read. A return value of 0 means no bytes can be read at this time. This may be because of an EOF, or because there is a granularity greater than one byte that the stream imposes on the returned data, and **size** is less than this granularity. (This will happen frequently for streams that need to return whole characters, because **Lstream\_read()** calls the reader function repeatedly until it has the number of bytes it wants or until 0 is returned.) The lstream functions do not treat a 0 return as EOF or do anything special; however, the calling function will interpret any 0 it gets back as EOF. This will normally not happen unless the caller calls **Lstream\_read()** with a very small size.

This function can be NULL if the stream is output-only.

**Bytecount writer** (Lstream \*stream, const unsigned char \*data, Bytecount size) [Lstream Method]

Send some data to the stream's end. Data to be sent is in **data** and is **size** bytes. Return the number of bytes sent. This function can send and return fewer bytes than is passed in; in that case, the function will just be called again until there is no data left or 0 is returned. A return value of 0 means that no more data can be currently stored, but there is no error; the data will be squirreled away until the writer can accept data. (This is useful, e.g., if you're dealing with a non-blocking file descriptor and are getting EWOULDBLOCK errors.) This function can be NULL if the stream is input-only.

**int rewinder** (Lstream \*stream) [Lstream Method]  
 Rewind the stream. If this is NULL, the stream is not seekable.

**int seekable\_p** (Lstream \*stream) [Lstream Method]  
 Indicate whether this stream is seekable—i.e. it can be rewound. This method is ignored if the stream does not have a rewind method. If this method is not present, the result is determined by whether a rewind method is present.

- int flusher** (Lstream \*stream) [Lstream Method]  
Perform any additional operations necessary to flush the data in this stream.
- int pseudo\_closer** (Lstream \*stream) [Lstream Method]
- int closer** (Lstream \*stream) [Lstream Method]  
Perform any additional operations necessary to close this stream down. May be `NULL`. This function is called when `Lstream_close()` is called or when the stream is garbage-collected. When this function is called, all pending data in the stream will already have been written out.
- Lisp\_Object marker** (Lisp\_Object lstream , void (\*markfun) (Lisp\_Object)) [Lstream Method]  
Mark this object for garbage collection. Same semantics as a standard `Lisp_Object` marker. This function can be `NULL`.

## 23 Consoles; Devices; Frames; Windows

### 23.1 Introduction to Consoles; Devices; Frames; Windows

A window-system window that you see on the screen is called a **frame** in Emacs terminology. Each frame is subdivided into one or more non-overlapping panes, called (confusingly) **windows**. Each window displays the text of a buffer in it. (See above on Buffers.) Note that buffers and windows are independent entities: Two or more windows can be displaying the same buffer (potentially in different locations), and a buffer can be displayed in no windows.

A single display screen that contains one or more frames is called a **display**. Under most circumstances, there is only one display. However, more than one display can exist, for example if you have a **multi-headed** console, i.e. one with a single keyboard but multiple displays. (Typically in such a situation, the various displays act like one large display, in that the mouse is only in one of them at a time, and moving the mouse off of one moves it into another.) In some cases, the different displays will have different characteristics, e.g. one color and one mono.

XEmacs can display frames on multiple displays. It can even deal simultaneously with frames on multiple keyboards (called **consoles** in XEmacs terminology). Here is one case where this might be useful: You are using XEmacs on your workstation at work, and leave it running. Then you go home and dial in on a TTY line, and you can use the already-running XEmacs process to display another frame on your local TTY.

Thus, there is a hierarchy console -> display -> frame -> window. There is a separate Lisp object type for each of these four concepts. Furthermore, there is logically a **selected console**, **selected display**, **selected frame** and **selected window**. Each of these objects is distinguished in various ways, such as being the default object for various functions that act on objects of that type. Note that every containing object remembers the “selected” object among the objects that it contains: e.g. not only is there a selected window, but every frame remembers the last window in it that was selected, and changing the selected frame causes the remembered window within it to become the selected window. Similar relationships apply for consoles to devices and devices to frames.

### 23.2 Point

Recall that every buffer has a current insertion position, called **point**. Now, two or more windows may be displaying the same buffer, and the text cursor in the two windows (i.e. **point**) can be in two different places. You may ask, how can that be, since each buffer has only one value of **point**? The answer is that each window also has a value of **point** that is squirreled away in it. There is only one selected window, and the value of “point” in that buffer corresponds to that window. When the selected window is changed from one window to another displaying the same buffer, the old value of **point** is stored into the old window’s “point” and the value of **point** from the new window is retrieved and made the value of **point** in the buffer. This means that **window-point** for the selected window is potentially inaccurate, and if you want to retrieve the correct value of **point** for a window, you must special-case on the selected window and retrieve the buffer’s point instead. This is related to why **save-window-excursion** does not save the selected window’s value of **point**.

### 23.3 Window Hierarchy

If a frame contains multiple windows (panes), they are always created by splitting an existing window along the horizontal or vertical axis. Terminology is a bit confusing here: to **split a window horizontally** means to create two side-by-side windows, i.e. to make a *vertical* cut in a window. Likewise, to **split a window vertically** means to create two windows, one above the other, by making a *horizontal* cut.

If you split a window and then split again along the same axis, you will end up with a number of panes all arranged along the same axis. The precise way in which the splits were made should not be important, and this is reflected internally. Internally, all windows are arranged in a tree, consisting of two types of windows, **combination** windows (which have children, and are covered completely by those children) and **leaf** windows, which have no children and are visible. Every combination window has two or more children, all arranged along the same axis. There are (logically) two subtypes of windows, depending on whether their children are horizontally or vertically arrayed. There is always one root window, which is either a leaf window (if the frame contains only one window) or a combination window (if the frame contains more than one window). In the latter case, the root window will have two or more children, either horizontally or vertically arrayed, and each of those children will be either a leaf window or another combination window.

Here are some rules:

1. Horizontal combination windows can never have children that are horizontal combination windows; same for vertical.
2. Only leaf windows can be split (obviously) and this splitting does one of two things: (a) turns the leaf window into a combination window and creates two new leaf children, or (b) turns the leaf window into one of the two new leaves and creates the other leaf. Rule (1) dictates which of these two outcomes happens.
3. Every combination window must have at least two children.
4. Leaf windows can never become combination windows. They can be deleted, however. If this results in a violation of (3), the parent combination window also gets deleted.
5. All functions that accept windows must be prepared to accept combination windows, and do something sane (e.g. signal an error if so). Combination windows *do* escape to the Lisp level.
6. All windows have three fields governing their contents: these are **hchild** (a list of horizontally-arrayed children), **vchild** (a list of vertically-arrayed children), and **bu er** (the buffer contained in a leaf window). Exactly one of these will be non-**nil**. Remember that **horizontally-arrayed** means “side-by-side” and **vertically-arrayed** means **one above the other**.
7. Leaf windows also have markers in their **start** (the first buffer position displayed in the window) and **pointm** (the window’s stashed value of **point**—see above) fields, while combination windows have **nil** in these fields.
8. The list of children for a window is threaded through the **next** and **prev** fields of each child window.
9. **Deleted windows can be undeleted** This happens as a result of restoring a window configuration, and is unlike frames, displays, and consoles, which, once deleted, can

never be restored. Deleting a window does nothing except set a special `dead` bit to 1 and clear out the `next`, `prev`, `hchild`, and `vchild` fields, for GC purposes.

- Most frames actually have two top-level windows—one for the minibuffer and one (the `root`) for everything else. The modeline (if present) separates these two. The `next` field of the root points to the minibuffer, and the `prev` field of the minibuffer points to the root. The other `next` and `prev` fields are `nil`, and the frame points to both of these windows. Minibuffer-less frames have no minibuffer window, and the `next` and `prev` of the root window are `nil`. Minibuffer-only frames have no root window, and the `next` of the minibuffer window is `nil` but the `prev` points to itself. (#### This is an artifact that should be fixed.)

## 23.4 The Window Object

Windows have the following accessible fields:

<code>frame</code>	The frame that this window is on.
<code>mini_p</code>	Non- <code>nil</code> if this window is a minibuffer window.
<code>buffer</code>	The buffer that the window is displaying. This may change often during the life of the window.
<code>dedicated</code>	Non- <code>nil</code> if this window is dedicated to its buffer.
<code>pointm</code>	This is the value of <code>point</code> in the current buffer when this window is selected; when it is not selected, it retains its previous value.
<code>start</code>	The position in the buffer that is the first character to be displayed in the window.
<code>force_start</code>	If this flag is non- <code>nil</code> , it says that the window has been scrolled explicitly by the Lisp program. This affects what the next redisplay does if <code>point</code> is off the screen: instead of scrolling the window to show the text around <code>point</code> , it moves <code>point</code> to a location that is on the screen.
<code>last_modified</code>	The <code>modified</code> field of the window's buffer, as of the last time a redisplay completed in this window.
<code>last_point</code>	The buffer's value of <code>point</code> , as of the last time a redisplay completed in this window.
<code>left</code>	This is the left-hand edge of the window, measured in columns. (The leftmost column on the screen is column 0.)
<code>top</code>	This is the top edge of the window, measured in lines. (The top line on the screen is line 0.)
<code>height</code>	The height of the window, measured in lines.
<code>width</code>	The width of the window, measured in columns.

- next** This is the window that is the next in the chain of siblings. It is `nil` in a window that is the rightmost or bottommost of a group of siblings.
- prev** This is the window that is the previous in the chain of siblings. It is `nil` in a window that is the leftmost or topmost of a group of siblings.
- parent** Internally, XEmacs arranges windows in a tree; each group of siblings has a parent window whose area includes all the siblings. This field points to a window's parent.
- Parent windows do not display buffers, and play little role in display except to shape their child windows. Emacs Lisp programs usually have no access to the parent windows; they operate on the windows at the leaves of the tree, which actually display buffers.
- hscroll** This is the number of columns that the display in the window is scrolled horizontally to the left. Normally, this is 0.
- use\_time** This is the last time that the window was selected. The function `get-lru-window` uses this field.
- display\_table**  
The window's display table, or `nil` if none is specified for it.
- update\_mode\_line**  
Non-`nil` means this window's mode line needs to be updated.
- base\_line\_number**  
The line number of a certain position in the buffer, or `nil`. This is used for displaying the line number of point in the mode line.
- base\_line\_pos**  
The position in the buffer for which the line number is known, or `nil` meaning none is known.
- region\_showing**  
If the region (or part of it) is highlighted in this window, this field holds the mark position that made one end of that region. Otherwise, this field is `nil`.

## 24 The Redisplay Mechanism

The redisplay mechanism is one of the most complicated sections of XEmacs, especially from a conceptual standpoint. This is doubly so because, unlike for the basic aspects of the Lisp interpreter, the computer science theories of how to efficiently handle redisplay are not well-developed.

When working with the redisplay mechanism, remember the Golden Rules of Redisplay:

1. It Is Better To Be Correct Than Fast.
2. Thou Shalt Not Run Elisp From Within Redisplay.
3. It Is Better To Be Fast Than Not To Be.

### 24.1 Critical Redisplay Sections

Within this section, we are defenseless and assume that the following cannot happen:

1. garbage collection
2. Lisp code evaluation
3. frame size changes

We ensure (3) by calling `hold_frame_size_changes()`, which will cause any pending frame size changes to get put on hold till after the end of the critical section. (1) follows automatically if (2) is met. `####` Unfortunately, there are some places where Lisp code can be called within this section. We need to remove them.

If `Fsignal()` is called during this critical section, we will `abort()`.

If garbage collection is called during this critical section, we simply return. `####` We should abort instead.

`####` If a frame-size change does occur we should probably actually be preempting redisplay.

### 24.2 Line Start Cache

The traditional scrolling code in Emacs breaks in a variable height world. It depends on the key assumption that the number of lines that can be displayed at any given time is fixed. This led to a complete separation of the scrolling code from the redisplay code. In order to fully support variable height lines, the scrolling code must actually be tightly integrated with redisplay. Only redisplay can determine how many lines will be displayed on a screen for any given starting point.

What is ideally wanted is a complete list of the starting buffer position for every possible display line of a buffer along with the height of that display line. Maintaining such a full list would be very expensive. We settle for having it include information for all areas which we happen to generate anyhow (i.e. the region currently being displayed) and for those areas we need to work with.

In order to ensure that the cache accurately represents what redisplay would actually show, it is necessary to invalidate it in many situations. If the buffer changes, the starting positions may no longer be correct. If a face or an extent has changed then the line heights may have altered. These events happen frequently enough that the cache can end up being

constantly disabled. With this potentially constant invalidation when is the cache ever useful?

Even if the cache is invalidated before every single usage, it is necessary. Scrolling often requires knowledge about display lines which are actually above or below the visible region. The cache provides a convenient light-weight method of storing this information for multiple display regions. This knowledge is necessary for the scrolling code to always obey the First Golden Rule of Redisplay.

If the cache already contains all of the information that the scrolling routines happen to need so that it doesn't have to go generate it, then we are able to obey the Third Golden Rule of Redisplay. The first thing we do to help out the cache is to always add the displayed region. This region had to be generated anyway, so the cache ends up getting the information basically for free. In those cases where a user is simply scrolling around viewing a buffer there is a high probability that this is sufficient to always provide the needed information. The second thing we can do is be smart about invalidating the cache.

TODO—Be smart about invalidating the cache. Potential places:

- Insertions at end-of-line which don't cause line-wraps do not alter the starting positions of any display lines. These types of buffer modifications should not invalidate the cache. This is actually a large optimization for redisplay speed as well.
- Buffer modifications frequently only affect the display of lines at and below where they occur. In these situations we should only invalidate the part of the cache starting at where the modification occurs.

In case you're wondering, the Second Golden Rule of Redisplay is not applicable.

## 24.3 Redisplay Piece by Piece

As you can begin to see redisplay is complex and also not well documented. Chuck no longer works on XEmacs so this section is my take on the workings of redisplay.

Redisplay happens in three phases:

1. Determine desired display in area that needs redisplay. Implemented by `redisplay.c`
2. Compare desired display with current display Implemented by `redisplay-output.c`
3. Output changes Implemented by `redisplay-output.c`, `redisplay-x.c`, `redisplay-msw.c` and `redisplay-tty.c`

Steps 1 and 2 are device-independent and relatively complex. Step 3 is mostly device-dependent.

Determining the desired display

Display attributes are stored in `display_line` structures. Each `display_line` consists of a set of `display_block`'s and each `display_block` contains a number of `rune`'s. Generally dynarr's of `display_line`'s are held by each window representing the current display and the desired display.

The `display_line` structures are tightly tied to buffers which presents a problem for redisplay as this connection is bogus for the modeline. Hence the `display_line` generation routines are duplicated for generating the modeline. This means that the modeline display code has many bugs that the standard redisplay code does not.

The guts of `display_line` generation are in `create_text_block`, which creates a single display line for the desired locale. This incrementally parses the characters on the current line and generates redisplay structures for each.

Gutter redisplay is different. Because the data to display is stored in a string we cannot use `create_text_block`. Instead we use `create_text_string_block` which performs the same function as `create_text_block` but for strings. Many of the complexities of `create_text_block` to do with cursor handling and selective display have been removed.



## 25 Extents

### 25.1 Introduction to Extents

Extents are regions over a buffer, with a start and an end position denoting the region of the buffer included in the extent. In addition, either end can be closed or open, meaning that the endpoint is or is not logically included in the extent. Insertion of a character at a closed endpoint causes the character to go inside the extent; insertion at an open endpoint causes the character to go outside.

Extent endpoints are stored using memory indices (see ‘`insdel.c`’), to minimize the amount of adjusting that needs to be done when characters are inserted or deleted.

(Formerly, extent endpoints at the gap could be either before or after the gap, depending on the open/closedness of the endpoint. The intent of this was to make it so that insertions would automatically go inside or out of extents as necessary with no further work needing to be done. It didn’t work out that way, however, and just ended up complexifying and buggifying all the rest of the code.)

### 25.2 Extent Ordering

Extents are compared using memory indices. There are two orderings for extents and both orders are kept current at all times. The normal or **display** order is as follows:

```
Extent A is ‘‘less than’’ extent B,
that is, earlier in the display order,
  if:    A-start < B-start,
  or if: A-start = B-start, and A-end > B-end
```

So if two extents begin at the same position, the larger of them is the earlier one in the display order (`EXTENT_LESS` is true).

For the e-order, the same thing holds:

```
Extent A is ‘‘less than’’ extent B in e-order,
that is, later in the buffer,
  if:    A-end < B-end,
  or if: A-end = B-end, and A-start > B-start
```

So if two extents end at the same position, the smaller of them is the earlier one in the e-order (`EXTENT_E_LESS` is true).

The display order and the e-order are complementary orders: any theorem about the display order also applies to the e-order if you swap all occurrences of “display order” and “e-order”, “less than” and “greater than”, and “extent start” and “extent end”.

### 25.3 Format of the Extent Info

An extent-info structure consists of a list of the buffer or string’s extents and a **stack of extents** that lists all of the extents over a particular position. The stack-of-extents info is used for optimization purposes—it basically caches some info that might be expensive to compute. Certain otherwise hard computations are easy given the stack of extents over a particular position, and if the stack of extents over a nearby position is known (because it

was calculated at some prior point in time), it's easy to move the stack of extents to the proper position.

Given that the stack of extents is an optimization, and given that it requires memory, a string's stack of extents is wiped out each time a garbage collection occurs. Therefore, any time you retrieve the stack of extents, it might not be there. If you need it to be there, use the `_force` version.

Similarly, a string may or may not have an `extent_info` structure. (Generally it won't if there haven't been any extents added to the string.) So use the `_force` version if you need the `extent_info` structure to be there.

A list of extents is maintained as a double gap array. One gap array is ordered by start index (the **display order**) and the other is ordered by end index (the **e-order**). Note that positions in an extent list should logically be conceived of as referring *to* a particular extent (as is the norm in programs) rather than sitting between two extents. Note also that callers of these functions should not be aware of the fact that the extent list is implemented as an array, except for the fact that positions are integers (this should be generalized to handle integers and linked list equally well).

A gap array is the same structure used by buffer text: an array of elements with a "gap" somewhere in the middle. Insertion and deletion happens by moving the gap to the insertion/deletion point, and then expanding/contracting as necessary. Gap arrays have a number of useful properties:

1. They are space efficient, as there is no need for next/previous pointers.
2. If the items in them are sorted, locating an item is fast –  $O(\log N)$ .
3. Insertion and deletion is very fast (constant time, essentially) if the gap is near (which favors localized operations, as will usually be the case). Even if not, it requires only a block move of memory, which is generally a highly optimized operation on modern processors.
4. Code to manipulate them is relatively simple to write.

An alternative would be balanced binary trees, which have guaranteed  $O(\log N)$  time for all operations (although the constant factors are not as good, and repeated localized operations will be slower than for a gap array). Such code is quite tricky to write, however.

## 25.4 Zero-Length Extents

Extents can be zero-length, and will end up that way if their endpoints are explicitly set that way or if their detachable property is `nil` and all the text in the extent is deleted. (The exception is open-open zero-length extents, which are barred from existing because there is no sensible way to define their properties. Deletion of the text in an open-open extent causes it to be converted into a closed-open extent.) Zero-length extents are primarily used to represent annotations, and behave as follows:

1. Insertion at the position of a zero-length extent expands the extent if both endpoints are closed; goes after the extent if it is closed-open; and goes before the extent if it is open-closed.
2. Deletion of a character on a side of a zero-length extent whose corresponding endpoint is closed causes the extent to be detached if it is detachable; if the extent is not

detachable or the corresponding endpoint is open, the extent remains in the buffer, moving as necessary.

Note that closed-open, non-detachable zero-length extents behave exactly like markers and that open-closed, non-detachable zero-length extents behave like the “point-type” marker in Mule.

## 25.5 Mathematics of Extent Ordering

The extents in a buffer are ordered by “display order” because that is that order that the redisplay mechanism needs to process them in. The e-order is an auxiliary ordering used to facilitate operations over extents. The operations that can be performed on the ordered list of extents in a buffer are

1. Locate where an extent would go if inserted into the list.
2. Insert an extent into the list.
3. Remove an extent from the list.
4. Map over all the extents that overlap a range.

(4) requires being able to determine the first and last extents that overlap a range.

NOTE: **overlap** is used as follows:

- two ranges overlap if they have at least one point in common. Whether the endpoints are open or closed makes a difference here.
- a point overlaps a range if the point is contained within the range; this is equivalent to treating a point  $P$  as the range  $[P, P]$ .
- In the case of an *extent* overlapping a point or range, the extent is normally treated as having closed endpoints. This applies consistently in the discussion of stacks of extents and such below. Note that this definition of overlap is not necessarily consistent with the extents that **map-extents** maps over, since **map-extents** sometimes pays attention to whether the endpoints of an extents are open or closed. But for our purposes, it greatly simplifies things to treat all extents as having closed endpoints.

First, define  $>$ ,  $<$ ,  $\leq$ , etc. as applied to extents to mean comparison according to the display order. Comparison between an extent  $E$  and an index  $I$  means comparison between  $E$  and the range  $[I, I]$ .

Also define  $e >$ ,  $e <$ ,  $e \leq$ , etc. to mean comparison according to the e-order.

For any range  $R$ , define  $R(0)$  to be the starting index of the range and  $R(1)$  to be the ending index of the range.

For any extent  $E$ , define  $E(next)$  to be the extent directly following  $E$ , and  $E(prev)$  to be the extent directly preceding  $E$ . Assume  $E(next)$  and  $E(prev)$  can be determined from  $E$  in constant time. (This is because we store the extent list as a doubly linked list.)

Similarly, define  $E(e - next)$  and  $E(e - prev)$  to be the extents directly following and preceding  $E$  in the e-order.

Now:

Let  $R$  be a range. Let  $F$  be the first extent overlapping  $R$ . Let  $L$  be the last extent overlapping  $R$ .

Theorem 1:  $R(1)$  lies between  $L$  and  $L(next)$ , i.e.  $L \leq R(1) < L(next)$ .

This follows easily from the definition of display order. The basic reason that this theorem applies is that the display order sorts by increasing starting index.

Therefore, we can determine  $L$  just by looking at where we would insert  $R(1)$  into the list, and if we know  $F$  and are moving forward over extents, we can easily determine when we've hit  $L$  by comparing the extent we're at to  $R(1)$ .

**Theorem 2:**  $F(e - prev) \ e < [1, R(0)] \ e \leq F$ .

This is the analog of Theorem 1, and applies because the e-order sorts by increasing ending index.

Therefore,  $F$  can be found in the same amount of time as operation (1), i.e. the time that it takes to locate where an extent would go if inserted into the e-order list. This is  $O(\log N)$ , since we are using gap arrays to manage extents.

Define a **stack of extents** (or **SOE**) as the set of extents (ordered in display order and e-order, just like for normal extent lists) that overlap an index  $I$ .

Now:

Let  $I$  be an index, let  $S$  be the stack of extents on  $I$  and let  $F$  be the first extent in  $S$ .

**Theorem 3:** The first extent in  $S$  is the first extent that overlaps any range  $[I, J]$ .

**Proof:** Any extent that overlaps  $[I, J]$  but does not include  $I$  must have a start index  $> I$ , and thus be greater than any extent in  $S$ .

Therefore, finding the first extent that overlaps a range  $R$  is the same as finding the first extent that overlaps  $R(0)$ .

**Theorem 4:** Let  $I2$  be an index such that  $I2 > I$ , and let  $F2$  be the first extent that overlaps  $I2$ . Then, either  $F2$  is in  $S$  or  $F2$  is greater than any extent in  $S$ .

**Proof:** If  $F2$  does not include  $I$  then its start index is greater than  $I$  and thus it is greater than any extent in  $S$ , including  $F$ . Otherwise,  $F2$  includes  $I$  and thus is in  $S$ , and thus  $F2 \geq F$ .

## 25.6 Extent Fragments

Imagine that the buffer is divided up into contiguous, non-overlapping **runs** of text such that no extent starts or ends within a run (extents that abut the run don't count).

An extent fragment is a structure that holds data about the run that contains a particular buffer position (if the buffer position is at the junction of two runs, the run after the position is used)—the beginning and end of the run, a list of all of the extents in that run, the **merged face** that results from merging all of the faces corresponding to those extents, the begin and end glyphs at the beginning of the run, etc. This is the information that redisplay needs in order to display this run.

Extent fragments have to be very quick to update to a new buffer position when moving linearly through the buffer. They rely on the stack-of-extents code, which does the heavy-duty algorithmic work of determining which extents overly a particular position.

## 26 Faces

Not yet documented.



## 27 Glyphs

Glyphs are graphical elements that can be displayed in XEmacs buffers or gutters. We use the term graphical element here in the broadest possible sense since glyphs can be as mundane as text or as arcane as a native tab widget.

In XEmacs, glyphs represent the uninstantiated state of graphical elements, i.e. they hold all the information necessary to produce an image on-screen but the image need not exist at this stage, and multiple screen images can be instantiated from a single glyph.

Glyphs are lazily instantiated by calling one of the glyph functions. This usually occurs within `redisplay` when `Fglyph_height` is called. Instantiation causes an image-instance to be created and cached. This cache is on a per-device basis for all glyphs except widget-glyphs, and on a per-window basis for widget-glyphs. The caching is done by `image_instantiate` and is necessary because it is generally possible to display an image-instance in multiple domains. For instance if we create a Pixmap, we can actually display this on multiple windows - even though we only need a single Pixmap instance to do this. If caching wasn't done then it would be necessary to create image-instances for every displayable occurrence of a glyph - and every usage - and this would be extremely memory and cpu intensive.

Widget-glyphs (a.k.a native widgets) are not cached in this way. This is because widget-glyph image-instances on screen are toolkit windows, and thus cannot be reused in multiple XEmacs domains. Thus widget-glyphs are cached on an XEmacs window basis.

Any action on a glyph first consults the cache before actually instantiating a widget.

### 27.1 Glyph Instantiation

Glyph instantiation is a hairy topic and requires some explanation. The guts of glyph instantiation is contained within `image_instantiate`. A glyph contains an image which is a specifier. When a glyph function - for instance `Fglyph_height` - asks for a property of the glyph that can only be determined from its instantiated state, then the glyph image is instantiated and an image instance created. The instantiation process is governed by the specifier code and goes through a series of steps:

- **Validation.** Instantiation of image instances happens dynamically - often within the guts of `redisplay`. Thus it is often not feasible to catch instantiator errors at instantiation time. Instead the instantiator is validated at the time it is added to the image specifier. This function is defined by `image_validate` and at a simple level validates keyword value pairs.
- **Duplication.** The specifier code by default takes a copy of the instantiator. This is reasonable for most specifiers but in the case of widget-glyphs can be problematic, since some of the properties in the instantiator - for instance callbacks - could cause infinite recursion in the copying process. Thus the image code defines a function - `image_copy_instantiator` - which will selectively copy values. This is controlled by the way that a keyword is defined either using `IIFORMAT_VALID_KEYWORD` or `IIFORMAT_VALID_NONCOPY_KEYWORD`. Note that the image caching and `redisplay` code relies on instantiator copying to ensure that current and new instantiators are actually different rather than referring to the same thing.

- Normalization. Once the instantiator has been copied it must be converted into a form that is viable at instantiation time. This can involve no changes at all, but typically involves things like converting file names to the actual data. This function is defined by `image_going_to_add` and `normalize_image_instantiator`.
- Instantiation. When an image instance is actually required for display it is instantiated using `image_instantiate`. This involves calling instantiate methods that are specific to the type of image being instantiated.

The final instantiation phase also involves a number of steps. In order to understand these we need to describe a number of concepts.

An image is instantiated in a **domain**, where a domain can be any one of a device, frame, window or image-instance. The domain gives the image-instance context and identity and properties that affect the appearance of the image-instance may be different for the same glyph instantiated in different domains. An example is the face used to display the image-instance.

Although an image is instantiated in a particular domain the instantiation domain is not necessarily the domain in which the image-instance is cached. For example a pixmap can be instantiated in a window but actually be cached on a per-device basis. The domain in which the image-instance is actually cached is called the **governing-domain**. A governing-domain is currently either a device or a window. Widget-glyphs and text-glyphs have a window as a governing-domain, all other image-instances have a device as the governing-domain. The governing domain for an image-instance is determined using the `governing-domain-image-instance` method.

## 27.2 Widget-Glyphs

### 27.3 Widget-Glyphs in the MS-Windows Environment

To Do

### 27.4 Widget-Glyphs in the X Environment

Widget-glyphs under X make heavy use of `lwlib` (see [Section 31.1 \[Lucid Widget Library\]](#), [page 171](#)) for manipulating the native toolkit objects. This is primarily so that different toolkits can be supported for widget-glyphs, just as they are supported for features such as menubars etc.

`Lwlib` is extremely poorly documented and quite hairy so here is my understanding of what goes on.

`Lwlib` maintains a set of `widget_instances` which mirror the hierarchical state of Xt widgets. I think this is so that widgets can be updated and manipulated generically by the `lwlib` library. For instance `update_one_widget_instance` can cope with multiple types of widget and multiple types of toolkit. Each element in the widget hierarchy is updated from its corresponding `widget_instance` by walking the `widget_instance` tree recursively.

This has desirable properties such as `lw_modify_all_widgets` which is called from `'glyphs-x.c'` and updates all the properties of a widget without having to know what the widget is or what toolkit it is from. Unfortunately this also has hairy properties such as

making the `lwlib` code quite complex. And of course `lwlib` has to know at some level what the widget is and how to set its properties.



## 28 Specifiers

Not yet documented.



## 29 Menus

A menu is set by setting the value of the variable `current-menubar` (which may be buffer-local) and then calling `set-menubar-dirty-flag` to signal a change. This will cause the menu to be redrawn at the next redisplay. The format of the data in `current-menubar` is described in ‘`menubar.c`’.

Internally the data in `current-menubar` is parsed into a tree of `widget_value`’s (defined in ‘`lwlib.h`’); this is accomplished by the recursive function `menu_item_descriptor_to_widget_value()`, called by `compute_menubar_data()`. Such a tree is deallocated using `free_widget_value()`.

`update_screen_menubars()` is one of the external entry points. This checks to see, for each screen, if that screen’s menubar needs to be updated. This is the case if

1. `set-menubar-dirty-flag` was called since the last redisplay. (This function sets the C variable `menubar_has_changed`.)
2. The buffer displayed in the screen has changed.
3. The screen has no menubar currently displayed.

`set_screen_menubar()` is called for each such screen. This function calls `compute_menubar_data()` to create the tree of `widget_value`’s, then calls `lw_create_widget()`, `lw_modify_all_widgets()`, and/or `lw_destroy_all_widgets()` to create the X-Toolkit widget associated with the menu.

`update_psheets()`, the other external entry point, actually changes the menus being displayed. It uses the widgets fixed by `update_screen_menubars()` and calls various X functions to ensure that the menus are displayed properly.

The menubar widget is set up so that `pre_activate_callback()` is called when the menu is first selected (i.e. mouse button goes down), and `menubar_selection_callback()` is called when an item is selected. `pre_activate_callback()` calls the function in `activate-menubar-hook`, which can change the menubar (this is described in ‘`menubar.c`’). If the menubar is changed, `set_screen_menubars()` is called. `menubar_selection_callback()` enqueues a menu event, putting in it a function to call (either `eval` or `call-interactively`) and its argument, which is the callback function or form given in the menu’s description.



## 30 Subprocesses

The fields of a process are:

<code>name</code>	A string, the name of the process.
<code>command</code>	A list containing the command arguments that were used to start this process.
<code>filter</code>	A function used to accept output from the process instead of a buffer, or <code>nil</code> .
<code>sentinel</code>	A function called whenever the process receives a signal, or <code>nil</code> .
<code>buffer</code>	The associated buffer of the process.
<code>pid</code>	An integer, the Unix process <code>id</code> .
<code>childp</code>	A flag, non- <code>nil</code> if this is really a child process. It is <code>nil</code> for a network connection.
<code>mark</code>	A marker indicating the position of the end of the last output from this process inserted into the buffer. This is often but not always the end of the buffer.
<code>kill_without_query</code>	If this is non- <code>nil</code> , killing XEmacs while this process is still running does not ask for confirmation about killing the process.
<code>raw_status_low</code>	
<code>raw_status_high</code>	These two fields record 16 bits each of the process status returned by the <code>wait</code> system call.
<code>status</code>	The process status, as <code>process-status</code> should return it.
<code>tick</code>	
<code>update_tick</code>	If these two fields are not equal, a change in the status of the process needs to be reported, either by running the sentinel or by inserting a message in the process buffer.
<code>pty_flag</code>	Non- <code>nil</code> if communication with the subprocess uses a <code>pty</code> ; <code>nil</code> if it uses a pipe.
<code>infd</code>	The file descriptor for input from the process.
<code>outfd</code>	The file descriptor for output to the process.
<code>subtty</code>	The file descriptor for the terminal that the subprocess is using. (On some systems, there is no need to record this, so the value is <code>-1</code> .)
<code>tty_name</code>	The name of the terminal that the subprocess is using, or <code>nil</code> if it is using pipes.



## 31 Interface to the X Window System

Mostly undocumented.

### 31.1 Lucid Widget Library

Lwlib is extremely poorly documented and quite hairy. The author(s) blame that on X, Xt, and Motif, with some justice, but also sufficient hypocrisy to avoid drawing the obvious conclusion about their own work.

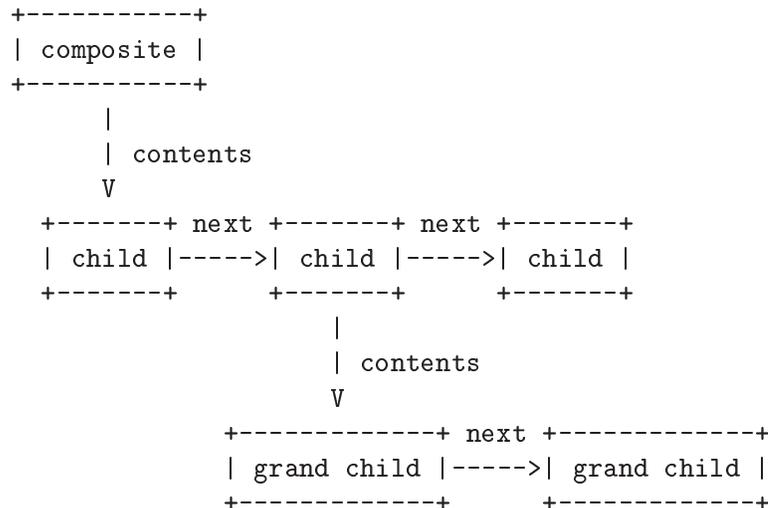
The Lucid Widget Library is composed of two more or less independent pieces. The first, as the name suggests, is a set of widgets. These widgets are intended to resemble and improve on widgets provided in the Motif toolkit but not in the Athena widgets, including menubars and scrollbars. Recent additions by Andy Piper integrate some “modern” widgets by Edward Falk, including checkboxes, radio buttons, progress gauges, and index tab controls (aka notebooks).

The second piece of the Lucid widget library is a generic interface to several toolkits for X (including Xt, the Athena widget set, and Motif, as well as the Lucid widgets themselves) so that core XEmacs code need not know which widget set has been used to build the graphical user interface.

#### 31.1.1 Generic Widget Interface

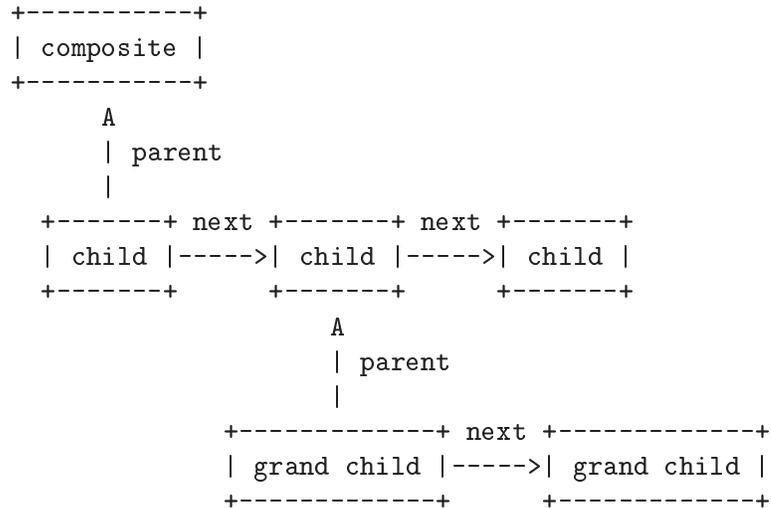
In general in any toolkit a widget may be a composite object. In Xt, all widgets have an X window that they manage, but typically a complex widget will have widget children, each of which manages a subwindow of the parent widget’s X window. These children may themselves be composite widgets. Thus a widget is actually a tree or hierarchy of widgets.

For each toolkit widget, lwlib maintains a tree of `widget_values` which mirror the hierarchical state of Xt widgets (including Motif, Athena, 3D Athena, and Falk’s widget sets). Each `widget_value` has `contents` member, which points to the head of a linked list of its children. The linked list of siblings is chained through the `next` member of `widget_value`.



The `widget_value` hierarchy of a composite widget with two simple children and one composite child.

The `widget_instance` structure maintains the inverse view of the tree. As for the `widget_value`, siblings are chained through the `next` member. However, rather than naming children, the `widget_instance` tree links to parents.



The `widget_value` hierarchy of a composite widget with two simple children and one composite child.

This permits widgets derived from different toolkits to be updated and manipulated generically by the `lwlib` library. For instance `update_one_widget_instance` can cope with multiple types of widget and multiple types of toolkit. Each element in the widget hierarchy is updated from its corresponding `widget_value` by walking the `widget_value` tree. This has desirable properties. For example, `lw_modify_all_widgets` is called from `'glyphs-x.c'` and updates all the properties of a widget without having to know what the widget is or what toolkit it is from. Unfortunately this also has its hairy properties; the `lwlib` code quite complex. And of course `lwlib` has to know at some level what the widget is and how to set its properties.

The `widget_instance` structure also contains a pointer to the root of its tree. Widget instances are further confi

### 31.1.2 Scrollbars

### 31.1.3 Menubars

### 31.1.4 Checkboxes and Radio Buttons

### 31.1.5 Progress Bars

### 31.1.6 Tab Controls

## Index

## A

allocation from frob blocks .....	95
allocation of objects in XEmacs Lisp .....	85
allocation, introduction to .....	85
allocation, low-level .....	98
Amdahl Corporation .....	4
Andressen, Marc .....	4
Assert .....	57
asynchronous subprocesses .....	77

## B

bars, progress .....	172
Baur, Steve .....	4, 5
Benson, Eric .....	3
binding; the specbinding stack; unwind-protects, dynamic .....	119
bindings, evaluation; stack frames; .....	117
bit vector .....	100
bridge, playing .....	9
Buchholz, Martin .....	4, 5
buffer lists .....	125
buffer object, the .....	127
buffer, the text in a .....	123
buffers and textual representation .....	123
buffers, introduction to .....	123
build-time dependencies .....	17
building, XEmacs from the perspective of .....	15
buttons, checkboxes and radio .....	172
byte positions, working with character and .....	47
Bytebpos .....	46
Bytecount .....	46
bytecount_to_charcount .....	48

## C

C code, rules when writing new .....	37
C vs. Lisp .....	11
callback routines, the event stream .....	115
caller-protects (GCPRole) .....	42
case table .....	75
catch and throw .....	119
CCL .....	137
character and byte positions, working with .....	47
character encoding, internal .....	136
character sets .....	133
character sets and encodings, Mule .....	133
character-related data types .....	46
characters, integers and .....	94
Charbpos .....	46
Charcount .....	46
charcount_to_bytecount .....	48
Charxpos .....	46
Check-Error .....	57

Check-Error-Message .....	57
Check-Message .....	57
checkboxes and radio buttons .....	172
closer .....	146
closure .....	23
code, an example of Mule-aware .....	52
code, general guidelines for writing Mule-aware .....	51
code, rules when writing new C .....	37
coding conventions .....	37
coding for Mule .....	45
coding rules, general .....	38
coding rules, naming .....	37
command builder, dispatching events; the .....	115
comments, writing good .....	43
Common Lisp .....	11
compact_string_chars .....	93
compiled function .....	101
compiler, the Lisp reader and .....	141
cons .....	99
conservative garbage collection .....	88
consoles; devices; frames; windows .....	147
consoles; devices; frames; windows, introduction to .....	147
control flow modules, editor-level .....	69
conversion to and from external data .....	48
converting events .....	115
copy-on-write .....	39
creating Lisp object types .....	56
critical redisplay sections .....	151
CVS techniques .....	61

## D

data dumping .....	106
data types, character-related .....	46
DEC_IBYTEPTR .....	48
dependencies, build-time .....	17
developers, techniques for XEmacs .....	53
devices; frames; windows, consoles; .....	147
devices; frames; windows, introduction to consoles; .....	147
Devin, Matthieu .....	3
dispatching events; the command builder .....	115
display order of extents .....	157
display-related Lisp objects, modules for other ..	72
displayable Lisp objects, modules for the basic ..	71
dumping .....	103
dumping address allocation .....	105
dumping and its justification, what is .....	103
dumping data descriptions .....	103
dumping object inventory .....	104
dumping overview .....	103
dumping phase .....	104
dumping, data .....	106

dumping, le loading .....	106
dumping, object relocation .....	107
dumping, pointers .....	106
dumping, putting back the pdump _opaques...	107
dumping, putting back the pdump _root_objects and pdump_weak_object_chains .....	107
dumping, putting back the pdump_root_struct_ptrs .....	107
dumping, reloading phase .....	106
dumping, remaining issues .....	107
dumping, reorganize the hash tables .....	107
dumping, the header .....	106
dynamic array .....	65
dynamic binding; the specbinding stack; unwind-protects .....	119
dynamic scoping .....	11
dynamic types .....	11
<b>E</b>	
editing operations, modules for standard .....	68
Emacs 19, GNU .....	3
Emacs 20, GNU .....	4
Emacs, a history of .....	1
encoding, internal character .....	136
encoding, internal string .....	136
encodings, internal Mule .....	135
encodings, Mule .....	134
encodings, Mule character sets and .....	133
Energize .....	2
Epoch .....	2, 4
error checking .....	53
EUC (Extended Unix Code), Japanese .....	134
evaluation .....	117
evaluation; stack frames; bindings .....	117
event gathering mechanism, specifics of the ...	110
event loop functions, other .....	115
event loop, events and the .....	109
event stream callback routines, the .....	115
event, specifics about the Lisp object .....	115
events and the event loop .....	109
events, converting .....	115
events, introduction to .....	109
events, main loop .....	109
events; the command builder, dispatching .....	115
Extbyte .....	47
Extended Unix Code, Japanese EUC .....	134
extent fragments .....	158
extent info, format of the .....	155
extent mathematics .....	157
extent ordering .....	155, 157
extents .....	155
extents, display order .....	157
extents, introduction to .....	155
extents, markers and .....	126
extents, zero-length .....	156
external data, conversion to and from .....	48
external widget .....	81
<b>F</b>	
faces .....	159
le system, modules for interfacing with the ...	74
flusher .....	146
fragments, extent .....	158
frames; windows, consoles; devices; .....	147
frames; windows, introduction to consoles; devices; .....	147
Free Software Foundation .....	1
frob blocks, allocation from .....	95
FSF .....	1
FSF Emacs .....	3, 4
function, compiled .....	101
<b>G</b>	
garbage collection .....	86
garbage collection - step by step .....	88
garbage collection protection .....	40, 86
garbage collection, conservative .....	88
garbage collection, invocation .....	88
garbage_collect_1 .....	89
gc_sweep .....	92
GCPRNG .....	86
global Lisp variables, adding .....	44
glyph instantiation .....	161
glyphs .....	161
GNU Emacs 19 .....	3
GNU Emacs 20 .....	4
Gosling, James .....	1, 11
Great Integral Type Renaming .....	31
Great Usenet Renaming .....	1
<b>H</b>	
Hackers (Steven Levy) .....	1
header files, inline functions .....	55
hierarchy of windows .....	148
history of Emacs, a .....	1
<b>I</b>	
lbyte .....	46
lbytes and lchars .....	126
lchar .....	46
lchars, lbytes and .....	126
Ignore-Ebola .....	58
Illinois, University of .....	4
INC_IBYTEPTR .....	48
inline functions .....	55
inline functions, headers .....	55
inside, XEmacs from the .....	19
instantiation, glyph .....	161
integers and characters .....	94
integral type renaming, great .....	31
interactive .....	69
interfacing with the le system, modules for ...	74

- interfacing with the operating system, modules for ..... 77
  - interfacing with X Windows, modules for ..... 80
  - internal character encoding ..... 136
  - internal Mule encodings ..... 135
  - internal string encoding ..... 136
  - internationalization, modules for ..... 81
  - interning ..... 27
  - interpreter and object system, modules for other aspects of the Lisp..... 75
  - itext\_lchar ..... 47
  - itext\_n\_addr..... 48
  - ITS (Incompatible Timesharing System) ..... 1
- J**
- Japanese EUC (Extended Unix Code) ..... 134
  - Java ..... 11
  - Java vs. Lisp ..... 11
  - JIS7 ..... 135
  - Jones, Kyle ..... 5
- K**
- Kaplan, Simon ..... 4
  - Known-Bug-Expect-Failure ..... 58
- L**
- Levy, Steven..... 1
  - library, Lucid Widget ..... 171
  - line start cache ..... 151
  - Lisp interpreter and object system, modules for other aspects of the..... 75
  - Lisp language, the ..... 11
  - Lisp modules, basic..... 65
  - Lisp object types, creating ..... 56
  - Lisp objects are represented in C, how..... 29
  - Lisp objects, allocation of in XEmacs ..... 85
  - Lisp objects, modules for other display-related .. 72
  - Lisp objects, modules for the basic displayable.. 71
  - Lisp primitives, writing ..... 40
  - Lisp reader and compiler, the ..... 141
  - Lisp vs. C..... 11
  - Lisp vs. Java..... 11
  - low-level allocation ..... 98
  - low-level modules..... 63
  - lrecords..... 95
  - lstream ..... 74
  - lstream functions ..... 144
  - lstream methods..... 145
  - lstream types..... 143
  - lstream, creating an..... 143
  - Lstream\_close ..... 145
  - Lstream\_fgetc ..... 144
  - Lstream\_flush ..... 144
  - Lstream\_fputc ..... 144
  - Lstream\_fungetc ..... 144
  - Lstream\_getc ..... 144
  - Lstream\_new..... 144
  - Lstream\_putc ..... 144
  - Lstream\_read ..... 144
  - Lstream\_reopen ..... 145
  - Lstream\_rewind ..... 145
  - Lstream\_set\_buffering ..... 144
  - Lstream\_ungetc ..... 144
  - Lstream\_unread ..... 145
  - Lstream\_write ..... 144
  - lstreams ..... 143
  - Lucid Emacs ..... 2
  - Lucid Inc. .... 2
  - Lucid Widget Library ..... 171
- M**
- macro hygiene..... 54
  - main loop ..... 109
  - major textual changes ..... 31
  - mark and sweep..... 86
  - mark method ..... 76, 97
  - mark\_object ..... 91
  - marker ..... 100
  - marker ..... 146
  - markers and extents ..... 126
  - matching ..... 128
  - mathematics of extent ordering ..... 157
  - MAX\_ICHAR\_LEN ..... 47
  - menubars..... 172
  - menus..... 167
  - merging a branch into the trunk ..... 61
  - merging attempts ..... 5
  - MIT ..... 1
  - Mlynarik, Richard ..... 4
  - modules for interfacing with the le system .... 74
  - modules for interfacing with the operating system ..... 77
  - modules for interfacing with X Windows ..... 80
  - modules for internationalization ..... 81
  - modules for other aspects of the Lisp interpreter and object system..... 75
  - modules for other display-related Lisp objects .. 72
  - modules for regression testing..... 82
  - modules for standard editing operations ..... 68
  - modules for the basic displayable Lisp objects.. 71
  - modules for the redisplay mechanism..... 73
  - modules, a summary of the various XEmacs ... 63
  - modules, basic Lisp..... 65
  - modules, editor-level control ow ..... 69
  - modules, low-level ..... 63
  - MS-Windows environment, widget-glyphs in the ..... 162
  - Mule character sets and encodings..... 133
  - Mule encodings..... 134
  - Mule encodings, internal ..... 135
  - MULE merged XEmacs appears ..... 5
  - Mule, coding for ..... 45

Mule-aware code, an example of.....	52	redisplay sections, critical .....	151
Mule-aware code, general guidelines for writing .....	51	regression testing, modules for.....	82
<b>N</b>			
NAS.....	79	reloading phase.....	106
native sound .....	79	relocating allocator .....	64
network connections .....	77	rename to XEmacs.....	5
network sound.....	79	renaming, integral types .....	31
Niksic, Hrvoje .....	5	renaming, text/char types .....	34
<b>O</b>			
obarrays.....	121	represented in C, how Lisp objects are.....	29
object system (abstractly speaking), the XEmacs .....	23	rewinder .....	145
object system, modules for other aspects of the Lisp interpreter and .....	75	RMS.....	1
object types, creating Lisp .....	56	<b>S</b>	
object, the buffer .....	127	scanner.....	76
object, the window .....	149	scoping, dynamic.....	11
objects are represented in C, how Lisp.....	29	scrollbars .....	172
objects in XEmacs Lisp, allocation of .....	85	searching.....	128
objects, modules for the basic displayable Lisp..	71	seekable_p .....	145
operating system, modules for interfacing with the .....	77	selections.....	80
outside, XEmacs from the .....	9	set_itext_ichar .....	47
<b>P</b>			
pane.....	72	Sexton, Harlan .....	3
permanent objects.....	26	Skip-Test-Unless .....	59
pi, calculating .....	9	skipped-test-reasons .....	59
point .....	147	sound, native.....	79
pointers dumping .....	106	sound, network .....	79
positions, working with character and byte .....	47	SPARCWorks .....	4
primitives, writing Lisp .....	40	specbinding stack; unwind-protects, dynamic binding; the .....	119
progress bars.....	172	special forms, simple.....	119
protection, garbage collection .....	86	specifiers.....	165
pseudo_closer .....	146	stack frames; bindings, evaluation; .....	117
Purify .....	53	Stallman, Richard .....	1
<b>Q</b>			
Quantify .....	53	string .....	100
<b>R</b>			
radio buttons, checkboxes and.....	172	string encoding, internal .....	136
read syntax .....	26	subprocesses .....	169
read-eval-print .....	9	subprocesses, asynchronous.....	77
reader .....	145	subprocesses, synchronous.....	77
reader and compiler, the Lisp.....	141	Sun Microsystems.....	4
reader's guide .....	37	sweep_bit_vectors_1 .....	94
redisplay mechanism, modules for the.....	73	sweep_lcrecords_1 .....	93
redisplay mechanism, the.....	151	sweep_strings .....	94
redisplay piece by piece.....	152	symbol.....	100
		symbol values.....	122
		symbols and variables.....	121
		symbols, introduction to .....	121
		synchronous subprocesses.....	77

**T**

tab controls .....	172
taxes, doing .....	9
techniques for XEmacs developers .....	53
TECO .....	1
temporary objects .....	26
test-emacs-test-file .....	57
testing, regression .....	57
text in a buffer, the .....	123
Text/Char Type Renaming .....	34
textual changes, major .....	31
textual representation, buffers and .....	123
Thompson, Chuck .....	4
throw, catch and .....	119
type renaming, integral .....	31
type renaming, text/char .....	34
types, dynamic .....	11
types, lstream .....	143
types, proper use of unsigned .....	45

**U**

University of Illinois .....	4
unsigned types, proper use of .....	45
unwind-protects, dynamic binding; the specbinding stack; .....	119

**V**

values, symbol .....	122
variables, adding global Lisp .....	44
variables, symbols and .....	121
vector .....	100
vector, bit .....	100
version 18, through .....	1
version 19, GNU Emacs .....	3
version 20, GNU Emacs .....	4

**W**

widget interface, generic .....	171
widget library, Lucid .....	171
widget-glyphs .....	162
widget-glyphs in the MS-Windows environment .....	162
widget-glyphs in the X environment .....	162
Win-Emacs .....	4
window (in Emacs) .....	72
window hierarchy .....	148
window object, the .....	149
window point internals .....	149
windows, consoles; devices; frames; .....	147
windows, introduction to consoles; devices; frames; .....	147
Wing, Ben .....	4
writer .....	145
writing good comments .....	43
writing Lisp primitives .....	40
writing Mule-aware code, general guidelines for .....	51
writing new C code, rules when .....	37

**X**

X environment, widget-glyphs in the .....	162
X Window System, interface to the .....	171
X Windows, modules for interfacing with .....	80
XEmacs .....	4
XEmacs from the inside .....	19
XEmacs from the outside .....	9
XEmacs from the perspective of building .....	15
XEmacs goes it alone .....	5
XEmacs object system (abstractly speaking), the .....	23

**Z**

Zawinski, Jamie .....	3
zero-length extents .....	156



## Short Contents

1	A History of Emacs .....	1
2	XEmacs From the Outside .....	9
3	The Lisp Language .....	11
4	XEmacs From the Perspective of Building .....	15
5	Build-Time Dependencies .....	17
6	XEmacs From the Inside .....	19
7	The XEmacs Object System (Abstractly Speaking) .....	23
8	How Lisp Objects Are Represented in C .....	29
9	Major Textual Changes .....	31
10	Rules When Writing New C Code .....	37
11	Regression Testing XEmacs .....	57
12	CVS Techniques .....	61
13	A Summary of the Various XEmacs Modules .....	63
14	Allocation of Objects in XEmacs Lisp .....	85
15	Dumping .....	103
16	Events and the Event Loop .....	109
17	Evaluation; Stack Frames; Bindings .....	117
18	Symbols and Variables .....	121
19	Buffers and Textual Representation .....	123
20	MULE Character Sets and Encodings .....	133
21	The Lisp Reader and Compiler .....	141
22	Lstreams .....	143
23	Consoles; Devices; Frames; Windows .....	147
24	The Redisplay Mechanism .....	151
25	Extents .....	155
26	Faces .....	159
27	Glyphs .....	161
28	Specifiers .....	165
29	Menus .....	167
30	Subprocesses .....	169
31	Interface to the X Window System .....	171
	Index .....	173



## Table of Contents

<b>1</b>	<b>A History of Emacs .....</b>	<b>1</b>
1.1	Through Version 18.....	1
1.2	Lucid Emacs .....	2
1.3	GNU Emacs 19.....	3
1.4	GNU Emacs 20.....	4
1.5	XEmacs .....	4
<b>2</b>	<b>XEmacs From the Outside .....</b>	<b>9</b>
<b>3</b>	<b>The Lisp Language .....</b>	<b>11</b>
<b>4</b>	<b>XEmacs From the Perspective of Building..</b>	<b>15</b>
<b>5</b>	<b>Build-Time Dependencies .....</b>	<b>17</b>
<b>6</b>	<b>XEmacs From the Inside .....</b>	<b>19</b>
<b>7</b>	<b>The XEmacs Object System (Abstractly Speaking) .....</b>	<b>23</b>
<b>8</b>	<b>How Lisp Objects Are Represented in C ...</b>	<b>29</b>
<b>9</b>	<b>Major Textual Changes .....</b>	<b>31</b>
9.1	Great Integral Type Renaming .....	31
9.2	Text/Char Type Renaming.....	34
<b>10</b>	<b>Rules When Writing New C Code .....</b>	<b>37</b>
10.1	A Reader's Guide to XEmacs Coding Conventions .....	37
10.2	General Coding Rules.....	38
10.3	Writing Lisp Primitives .....	40
10.4	Writing Good Comments.....	43
10.5	Adding Global Lisp Variables.....	44
10.6	Proper Use of Unsigned Types.....	45
10.7	Coding for Mule.....	45
10.7.1	Character-Related Data Types .....	46
10.7.2	Working With Character and Byte Positions .....	47
10.7.3	Conversion to and from External Data.....	48
10.7.4	General Guidelines for Writing Mule-Aware Code....	51
10.7.5	An Example of Mule-Aware Code .....	52
10.7.6	Mule-izing Code .....	53
10.8	Techniques for XEmacs Developers .....	53

<b>11</b>	<b>Regression Testing XEmacs .....</b>	<b>57</b>
<b>12</b>	<b>CVS Techniques .....</b>	<b>61</b>
12.1	Merging a Branch into the Trunk .....	61
<b>13</b>	<b>A Summary of the Various XEmacs Modules</b> <b>.....</b>	<b>63</b>
13.1	Low-Level Modules .....	63
13.2	Basic Lisp Modules .....	65
13.3	Modules for Standard Editing Operations .....	68
13.4	Editor-Level Control Flow Modules .....	69
13.5	Modules for the Basic Displayable Lisp Objects .....	71
13.6	Modules for other Display-Related Lisp Objects .....	72
13.7	Modules for the Redisplay Mechanism .....	73
13.8	Modules for Interfacing with the File System .....	74
13.9	Modules for Other Aspects of the Lisp Interpreter and Object System .....	75
13.10	Modules for Interfacing with the Operating System .....	77
13.11	Modules for Interfacing with X Windows .....	80
13.12	Modules for Internationalization .....	81
13.13	Modules for Regression Testing .....	82
<b>14</b>	<b>Allocation of Objects in XEmacs Lisp .....</b>	<b>85</b>
14.1	Introduction to Allocation .....	85
14.2	Garbage Collection .....	86
14.3	GCPR0ing .....	86
14.4	Garbage Collection - Step by Step .....	88
14.4.1	Invocation .....	88
14.4.2	garbage_collect_1 .....	89
14.4.3	mark_object .....	91
14.4.4	gc_sweep .....	92
14.4.5	sweep_lrecords_1 .....	93
14.4.6	compact_string_chars .....	93
14.4.7	sweep_strings .....	94
14.4.8	sweep_bit_vectors_1 .....	94
14.5	Integers and Characters .....	94
14.6	Allocation from Frob Blocks .....	95
14.7	lrecords .....	95
14.8	Low-level allocation .....	98
14.9	Cons .....	99
14.10	Vector .....	100
14.11	Bit Vector .....	100
14.12	Symbol .....	100
14.13	Marker .....	100
14.14	String .....	100
14.15	Compiled Function .....	101

<b>15</b>	<b>Dumping.....</b>	<b>103</b>
15.1	What is dumping and its justification .....	103
15.2	Overview .....	103
15.3	Data descriptions.....	103
15.4	Dumping phase .....	104
15.4.1	Object inventory .....	104
15.4.2	Address allocation.....	105
15.4.3	The header .....	106
15.4.4	Data dumping .....	106
15.4.5	Pointers dumping .....	106
15.5	Reloading phase.....	106
15.5.1	File loading.....	106
15.5.2	Putting back the pdump_opaques .....	107
15.5.3	Putting back the pdump_root_struct_ptrs .....	107
15.5.4	Object relocation .....	107
15.5.5	Putting back the pdump_root_objects and pdump_weak_object_chains.....	107
15.5.6	Reorganize the hash tables.....	107
15.6	Remaining issues .....	107
<b>16</b>	<b>Events and the Event Loop.....</b>	<b>109</b>
16.1	Introduction to Events.....	109
16.2	Main Loop.....	109
16.3	Specifics of the Event Gathering Mechanism .....	110
16.4	Specifics About the Emacs Event .....	115
16.5	The Event Stream Callback Routines .....	115
16.6	Other Event Loop Functions .....	115
16.7	Converting Events.....	115
16.8	Dispatching Events; The Command Builder .....	115
<b>17</b>	<b>Evaluation; Stack Frames; Bindings.....</b>	<b>117</b>
17.1	Evaluation .....	117
17.2	Dynamic Binding; The specbinding Stack; Unwind-Protects .....	119
17.3	Simple Special Forms .....	119
17.4	Catch and Throw.....	119
<b>18</b>	<b>Symbols and Variables .....</b>	<b>121</b>
18.1	Introduction to Symbols .....	121
18.2	Obarrays .....	121
18.3	Symbol Values .....	122

<b>19</b>	<b>Buffers and Textual Representation.....</b>	<b>123</b>
19.1	Introduction to Buffers .....	123
19.2	The Text in a Buffer .....	123
19.3	Buffer Lists .....	125
19.4	Markers and Extents .....	126
19.5	Ibytes and Ichars .....	126
19.6	The Buffer Object .....	126
19.7	Searching and Matching .....	128
<b>20</b>	<b>MULE Character Sets and Encodings....</b>	<b>133</b>
20.1	Character Sets .....	133
20.2	Encodings .....	134
20.2.1	Japanese EUC (Extended Unix Code) .....	134
20.2.2	JIS7 .....	135
20.3	Internal Mule Encodings .....	135
20.3.1	Internal String Encoding .....	136
20.3.2	Internal Character Encoding .....	136
20.4	CCL .....	137
<b>21</b>	<b>The Lisp Reader and Compiler .....</b>	<b>141</b>
<b>22</b>	<b>Lstreams .....</b>	<b>143</b>
22.1	Creating an Lstream .....	143
22.2	Lstream Types .....	143
22.3	Lstream Functions .....	144
22.4	Lstream Methods .....	145
<b>23</b>	<b>Consoles; Devices; Frames; Windows.....</b>	<b>147</b>
23.1	Introduction to Consoles; Devices; Frames; Windows .....	147
23.2	Point .....	147
23.3	Window Hierarchy .....	148
23.4	The Window Object .....	149
<b>24</b>	<b>The Redisplay Mechanism .....</b>	<b>151</b>
24.1	Critical Redisplay Sections .....	151
24.2	Line Start Cache .....	151
24.3	Redisplay Piece by Piece .....	152
<b>25</b>	<b>Extents .....</b>	<b>155</b>
25.1	Introduction to Extents .....	155
25.2	Extent Ordering .....	155
25.3	Format of the Extent Info .....	155
25.4	Zero-Length Extents .....	156
25.5	Mathematics of Extent Ordering .....	157
25.6	Extent Fragments .....	158

<b>26</b>	<b>Faces .....</b>	<b>159</b>
<b>27</b>	<b>Glyphs.....</b>	<b>161</b>
27.1	Glyph Instantiation.....	161
27.2	Widget-Glyphs.....	162
27.3	Widget-Glyphs in the MS-Windows Environment .....	162
27.4	Widget-Glyphs in the X Environment.....	162
<b>28</b>	<b>Specifiers .....</b>	<b>165</b>
<b>29</b>	<b>Menus .....</b>	<b>167</b>
<b>30</b>	<b>Subprocesses .....</b>	<b>169</b>
<b>31</b>	<b>Interface to the X Window System .....</b>	<b>171</b>
31.1	Lucid Widget Library.....	171
31.1.1	Generic Widget Interface.....	171
31.1.2	Scrollbars.....	172
31.1.3	Menubars.....	172
31.1.4	Checkboxes and Radio Buttons .....	172
31.1.5	Progress Bars.....	172
31.1.6	Tab Controls .....	172
	<b>Index .....</b>	<b>173</b>

