

The UNIX Programming Environment

This handout was compiled by many other past CS107 and CS193D lecturers and TAs.

Introduction

In writing C and C++ programs to run under Unix, there are several concepts and tools that turn out to be quite useful. The most obvious difference, if you are coming from a PC or Macintosh programming background, is that the tools are separate entities, not components in a tightly coupled environment like Metrowerks CodeWarrior or Microsoft Visual C++. The appendix at the end of the handout gives a summary of some basic UNIX and EMACS commands.

The most important tools in this domain are the editor, the compiler, the linker, the `make` utility, and the debugger. There are a variety of choices as far as "which compiler" or "which editor", but the choice is usually one of personal preference. The choice of editor, however, is almost a religious issue. EMACS integrates well with the other tools, has a nice graphical interface, and is almost an operating system unto itself, so we will encourage its use.

Caveat

This handout is not meant to give a comprehensive discussion of any of the tools discussed, rather it should be used as an introduction to getting started with the tools. If you find that you need more information, all of the programs have extensive man pages and `xinfo` entries, and `gdb` has some on-line help for all of its commands. These man pages list numerous bits of information, some of which will appear to be quite attractive, but if you do not know or understand what an option does, especially for the compiler, please do not use it just because it sounds nice.

Also, O'Reilly & Associates publishes a pretty good set of references for these basic tools, the titles of which resemble "UNIX in a Nutshell". These are not required or endorsed (for the record), but may be useful if you get lost.

The Editing Process

The first frontier a person is confronted with when programming is the editing environment. As mentioned, the choice of editor in UNIX was and still is, for various valid and not so valid reasons, pretty much set in stone as the 11th commandment: Whosoever writes programs with EMACS¹ shall live without bugs, or else... EMACS has pretty much all the features you will possibly need in an editor, and then some; the only problem is finding out how to invoke them as needed. This section is only meant as a quick reference to get you started. To start editing a new or existing file using EMACS, simply type this to the UNIX prompt:

`emacs filename`

where *filename* is the file to be edited. Once EMACS has started up, text entry works much the same way as in any other text editors, i.e. the default mode is INSERT mode where text being

¹ In the 4th Century AD, Rome was plagued by the worshipers of the Church of Emacs. They never bathed. They would gather in large, foul-smelling groups and harass small children.... - Cornell Quote

typed is inserted at the current cursor position. All the fancy editing commands, such as find-and-replace, are invoked through typing special key sequences (there really isn't a great mouse-driven UNIX editor for writing code). Two important key sequences to remember are: `^x` (holding down the "ctrl" key while typing 'x') and `[esc]-x` (simply pressing the "esc" key followed by typing 'x'), both of which are used to *start* command sequences. Note that in most user manuals for EMACS, the "esc" key is actually referred to as the "Meta" key (let's not bother explaining that one). Therefore, you will often see commands prefaced with the key sequence `M-x`, as opposed to the `[esc]-x` that we will use in this handout. Since the two are pretty much synonymous, we'll stick with explicitly using the "esc" key.

Now let's see some examples of these mysterious "command sequences". For instance, to save the file being edited the sequence is `^x^s`. To exit (and be prompted to save) EMACS, the sequence is `^x^c`. To open another file within EMACS, the sequence is `^x^f` ('f' for "find file"). This sequence can be used to open an existing file as well as a new file. If you have multiple files open, EMACS stores them in different "buffers." To switch from one buffer to another (very handy when you are editing a `.c` source file and need to refer to the prototypes and definitions in the `.h` header file), you use the key sequence `^x` followed by typing 'b' (without holding down "ctrl" when typing 'b'). You can then enter the name of the file to switch to the corresponding buffer (a default name is provided for fast switching).

The arrow keys usually work as the cursor movement keys, but in the event that they don't, `^f` is for moving forward (right), `^b` is for moving backward (left), `^p` is for moving to the previous line (up), and `^n` is for moving to the next line (down). `[esc]-<` ([esc] followed by '<') moves to the beginning of the file, and `[esc]->` moves to the end of the file. Finally, `^v` does a "page down" and `[esc]-v` does a "page up." **Note that the [esc] key is simply pressed and not held, whereas the [ctrl] key is held in all cases.** For a listing of some other functions, please consult the appendix at the end of this handout.

Quite frequently you'll make typing mistakes. The backspace key usually deletes the character before the text cursor, but sometimes it is strangely bound to invoke help! In case that fails, try the delete key (the one on the mini-keypad that also contains page up, page down, etc). You can also delete a whole line at a time by using `^k`.

Another important editing feature is copy/cut and paste. To copy or cut, you first have to select a region of text. To tell `emacs` where this region begins, use `[esc]-@` or `^[space]` (control and space simultaneously). Then move the text cursor to the end of the region. Alternatively, if you're running `emacs` under X Windows, you can use your mouse by pressing down on the left button at the start of the region and "drag" your mouse cursor to the end. To copy the region, use `[esc]-w`, while to cut the region, use `^w`. To paste, use `^y`.

For a listing of some other functions, please consult the appendix at the end of this handout.


A variant of `emacs` is `xemacs`, which is essentially the same but with a more modern-looking user interface (buttons, menus, dialogs for opening files, etc). Nearly all commands that `emacs` accept are valid in `xemacs`.

The Compilation Process

Before going into detail about the actual tools themselves, it is useful to review what happens during the construction of an executable program. There are actually two phases in the process, *compilation* and *linking*. The individual source files must first be *compiled* into object modules. These object modules contain a system dependent, relocatable representation of the program as described in the source file. The individual object modules are then *linked* together to produce a single executable file which the system loader can use when the program is actually invoked. (This process is illustrated by the diagram on the next page.) These phases are often combined with the `gcc` command, but it is quite useful to separate them when using `make`.

For example, the command: `gcc -o prog main.c module1.c module2.c` can be broken down into the four steps of:

```
gcc -c main.c
gcc -c module1.c
gcc -c module2.c
gcc -o prog main.o module1.o module2.o
```



Compiler/Linker

Although there are a lot of different compilers out there, we “guarantee” that all of the problems can be solved using the GNU C Compiler, `gcc` to its friends. However, because C++ requires additional libraries, you should use `g++`, which is just a wrap-around of `gcc`, to compile C++ programs. Using `g++` has several advantages, it is pretty much ANSI compliant, available on a variety of different platforms and, more importantly for you, it works pretty reliably. The current version of `g++` installed on the L&IR machines is 2.8.1, and directly compiles C, C++, and Objective-C.

Running g++

Even though it is called a compiler, `g++` is used as both a compiler and linker. The general form for invoking `g++` is :

```
g++ <option flags> <file list>
```

where `<option flags>` is a list of command flags that control how the compiler works, and `<file list>` is a list of files, source or object, that `g++` is being directed to process. It is not, however, commonly invoked directly from the command line, that is what `makefiles` are for. If

g++ is unable to process the files correctly it will print error messages on standard error. Some of these error messages, however, will be caused by g++ trying to recover from a previous error so it is best to try to tackle the errors in order.

Command-line options

Like almost all UNIX programs g++ has myriad options that control almost every aspect of its actions. However, most of these options deal with system dependent features and do not concern us. The most useful option flags for us are: `-c`, `-o`, `-g`, `-Wall`, `-I`, `-L`, and `-l`.

- `-c` Requests that g++ compile the specific source file directly into an object file without going through the linking stage. This is important for compiling only those files that have changed rather than the whole project.
- `-o file` Specifies that you want the compiler's output to be named *file*. If this option is not specified, the default is to create a file 'a.out' if you are linking object files into an executable. If you are compiling a source file (with suffix .c for files written in C) into an object file, the default name for the object file is simply the original name with the '.c' replaced with '.o'. This option is generally only used for creating an application with a specific name (during linking), rather than for making the names of object files differ from the default *source-filename.o*.
- `-g` Directs the compiler to produce debugging information. We recommend that you always compile your source with this option set, since we encourage you to gain proficiency using the debugger (we recommend gdb).

Note – The debugging information generated is for gdb, and could possibly cause problems with dbx. This is because there is typically more information stored for gdb that dbx will barf on. Additionally, on some systems, some MIPS based machines for example, this information cannot encode full symbol information and some debugger features may be unavailable.

- `-Wall` Give warnings about a lot of syntactically correct but dubious constructs. Think of this option as being a way to do a simple form of style checking. Again, we highly recommend that you compile your code with this option set.

Most of the time the constructs that are flagged are actually incorrect usages, but there are occasionally instances where they are what you really want. Instead of simply ignoring these warnings there are simple workarounds for almost all of the warnings if you insist on doing things this way.

This sort of contrived snippet is a commonly used construct in C to set and test a variable in as few lines as possible :

```
int flag;

if (flag = IsPrime(13)) {
    ...
}
```

The compiler will give a warning about a possibly unintended assignment. This is because it is more common to have a boolean test in the `if` clause using the

equality operator `==` rather than to take advantage of the return value of the assignment operator. This snippet could better be written as :

```
int flag;

if ((flag = IsPrime(13)) != 0) {
    ...
}
```

so that the test for the 0 value is made explicit. The code generated will be the same, and it will make us and the compiler happy at the same time. Alternately, you can enclose the entire test in another set of parentheses to indicate your intentions.

-I*dir* Adds the directory *dir* to the list of directories searched for include files. This will be important for any additional files that we give you. There are a variety of standard directories that will be searched by the compiler by default, for standard library and system header files, but since we do not have root access we cannot just add our files to these locations.

There is no space between the option flag and the directory name.

-l*lib* Search the library named *lib* for unresolved names when linking. The actual name of the file will be *liblib.a*, and must be found in either the default locations for libraries or in a directory added with the `-L` flag.

The position of the `-l` flag in the option list is important because the linker will not go back to previously examined libraries to look for unresolved names. For example, if you are using a library that requires the math library it must appear before the math library on the command line otherwise a link error will be reported.

Again, there is no space between the option flag and the library file name.

-L*dir* Adds the directory *dir* to the list of directories searched for library files specified by the `-l` flag. Here too, there is no space between the option flag and the library directory name.

The make utility

As many of you probably already know, typing the entire command line to compile a program turns out to be a somewhat complicated and tedious affair. What the `make` utility does is to allow the programmer to write out a specification of all of the modules that go into creating an application, and how these modules need to be assembled to create the program. The `make` facility manages the execution of the necessary build commands (compiling, linking, loading etc.). In doing so, it also recognizes that only those files which have been changed need be rebuilt. Thus a properly constructed `makefile` can save a great deal of compilation time. Some people are “afraid” of `make` and its corresponding `makefile`, but in actuality creating a `makefile` is a pretty simple affair.

Running make

Invoking the `make` program is really simple, just type `make` at the shell prompt, or if you are an EMACS aficionado `M-x compile` will do basically the same thingⁱⁱ. Either of these commands will cause `make` to look in the current directory for a file called ‘`Makefile`’ or ‘`makefile`’ for the build instructions. If there is a problem building one of the targets along the way the error messages will appear on standard error or the EMACS ‘`compilation`’ buffer if you invoked `make` from within EMACS.

Makefile-craft

A `makefile` consists of a series of `make` variable definitions and dependency rules. A variable in a `makefile` is a name defined to represent some string of text. This works much like macro replacement in the C compiler’s pre-processor. Variables are most often used to represent a list of directories to search, options for the compiler, and names of programs to run. A variable is “declared” when it is set to a value. For example, the line :

```
CC = g++
```

will create a variable named `CC`, and set its value to be `gcc`. The name of the variable is case sensitive, and traditionally `make` variable names are in all capital letters.

While it is possible to define your own variables there are some that are considered ‘standard,’ and using them along with the default rules makes writing a `makefile` much easier. For the purposes of this class the important variables are: `CC`, `CFLAGS`, and `LDFLAGS`.

CC The name of the C compiler, this will default to `cc` in most versions of `make`. Please make sure that you set this to be `gcc` or `g++` since `cc` is not ANSI compliant on the LaIR SparcStations.

CFLAGS A list of options to pass on to the C compiler for all of your source files. This is commonly used to set the include path to include non-standard directories or build debugging versions, the `-I` and `-g` compiler flags. Sometimes this is called `CPPFLAGS` instead for C++ programs.

ⁱ ‘M-x’ means hold the ‘meta’ key down while hitting the ‘x’ key. If your keyboard does not have a ‘meta’ key then the ‘ESC’ will do the same thing. Hit the ‘ESC’ key and then the ‘x’ key. Do not hold down the ‘ESC’ key or else it will put you into eval mode.

LDFLAGS A list of options to pass on to the linker. This is most commonly used to set the library search path to non-standard directories and to include application specific library files, the `-L` and `-l` compiler flags.

Referencing the value of a variable is done by having a `$` followed by the name of the variable within parenthesis or curly braces. For example:

```
CFLAGS = -g -I/usr/class/cs193d/include
$(CC) $(CFLAGS) -c binky.c
```

The first line sets the value of the variable `CFLAGS` to turn on debugging information and add the directory `/usr/class/cs193d/include` to the include file search path. The second line uses the value of the variable `CC` as the name of the compiler to use passing to it the compiler options set in the previous line. If you use a variable that has not been previously set in the makefile, `make` will use the empty definition, an empty string.

The second major component of makefiles are dependency/build rules. A rule tells how to make a target based on changes to a list of certain files. The ordering of the rules in the makefile does not make any difference, except that the first rule is considered to be the default rule. The default rule is the rule that will be invoked when `make` is called without arguments, the usual way. If, however, you know exactly which rule you want to invoke you can name it directly with an argument to `make`. For example, if my makefile had a rule for `clean`, the command line `make clean` would invoke the actions listed after the `clean` label, more on actions later.

A rule generally consists of two lines, a dependency list and a command list. Here is an example:

```
binky.o : binky.c binky.h akbar.h
<tab>$(CC) $(CFLAGS) -c binky.c
```

The first line says that the object file `binky.o` must be rebuilt whenever any of `binky.c`, `binky.h`, or `akbar.h` are changed. The target `binky.o` is said to depend on these three files. Basically, an object file depends on its source file and any non-system files that it includes.

The second lineⁱⁱⁱ lists the commands that must be taken in order to rebuild `binky.o`, invoking the C compiler with whatever compiler options have been previously set. These lines must be indented with a `<tab>` character, just using spaces will not work! Because of this, make sure you are not in a mode which might substitute space characters for an actual tab. In particular, the "indented-text-mode" always "tabs" to the same indent level as the previous line, using spaces if the previous line were indented less than the standard full 8 spaces. This is also a problem when using copy/paste from some terminal programs. To check whether you have a tab character on that line, move to the beginning of that line and try to move "right" (`^f`). If the cursor skips 8 spaces to the right, you have a tab. If it moves space by space, then you need to delete the spaces and retype a tab character.

For "standard" compilations^{iv}, the second line can be omitted, and `make` will use the default build rule for the source file based on its extension, `.c` for C files. The default build rule that `make` uses for C files looks like this :

```
$(CC) $(CFLAGS) -c <source-file>
```

ⁱⁱⁱ The second line can actually be more than one line if multiple commands need to be done for a single target. In this class, however, we will not be doing anything that requires multiple commands per target.

^{iv} Most versions of `make` handle at least FORTRAN, C, and C++.

Here is a “complete” makefile for your reading pleasure.

```
CC = g++c
CFLAGS = -g -I/usr/class/cs193d/include
LDFLAGS = -L/usr/class/cs193d/lib -lgraph

PROG = example
HDRS = binky.h akbar.h defs.h
SRCS = main.c binky.cc akbar.cc
OBJS = main.o binky.o akbar.o

$(PROG) : $(OBJECTS)
    $(CC) -o $(PROG) $(LDFLAGS) $(OBJS)

clean :
    rm -f core $(PROG) $(OBJS)

TAGS : $(SRCS) $(HDRS)
    etags -t $(SRCS) $(HDRS)

main.o : binky.h akbar.h defs.h
binky.o : binky.h
akbar.o : akbar.h defs.h
```

This makefile includes two extra targets, in addition to building the executable: `clean` and `TAGS`. These are commonly included in makefiles to make your life as a programmer a little bit easier. The `clean` target is used to remove all of the object files and the executable so that you can start the build process from scratch^v, you will need to do this if you move to a system with a different architecture from where your object libraries were originally compiled, since source code is compiled differently on different types of machines. The `TAGS` rule creates a tag file that most Unix editors can use to search for symbol definitions^{vi}.

Compiling in Emacs

The Emacs editor provides support for the compile process. To compile your code from Emacs, type `M-x compile`. You will be prompted for a compile command. If you have a makefile, just type `make` and hit return. The makefile will be read and the appropriate commands executed. The Emacs buffer will split at this point, and compile errors will be brought up in the newly created buffer. In order to go to the line where a compile error occurred, place the cursor on the line which contains the error message and hit `^c^c`. This will jump the cursor to the line in your code where the error occurred.

^v It also removes any ‘core’ files that you might have lying around, not that there should be any.

^{vi} Use ‘`M-x find-tag`’ or ‘`M-.`’ in emacs to search for a symbol within emacs. Tags are a useful thing.

The Debugger (gdb)

During the course of the quarter you may run into a bug or two in your programs^{vii}. There are a variety of different techniques for finding these “anomalies,” but a good debugger can make the job a lot easier and faster. We recommend the GNU debugger, since it basically stomps on `dbx` in every possible area and works nicely with the `gcc` compiler we recommend. Other nice debugging environments include `ups` and `CodeCenter`, but these are not as universally available as `gdb`, and in the case of `CodeCenter` not as cheaply. While `gdb` does not have a flashy graphical interface as do the others, it is a powerful tool that provides the knowledgeable programmer with all of the information she could possibly want and then some.

This section does not come anywhere close to describing all of the features of `gdb`, but will hit on the high points. There is on-line help for `gdb` which can be seen by using the `help` command from within `gdb`. If you want more information try `xinfo` if you are logged onto the console of a machine with an X display or use the info-browser mode from within `emacs`.

A debugger is invaluable to a programmer because it eases the process of discovering and repairing bugs at run-time. In most programs of any significant size, it is not possible to determine all of the bugs in a program at compile-time because of oversights and misconceptions about the problem that the application is designed to solve.

The way debuggers allow you to find bugs is by allowing you to run your program on a line-by-line basis, pausing the program at times or conditions that you specify and allowing you to examine variables, registers, the run time stack and other facets of program state while paused.

Sometimes these bugs result in program crashes (a.k.a. "core dumps", "register dumps", etc.) that bring your program to a halt with a message like "Segmentation Violation" or the like. If your program has a severe bug that causes a program crash, the debugger will "catch" the signal sent by the processor that indicates the error it found, and allow you to further examine the program. This information can be quite valuable when trying to reason about what caused your program to die, all segmentation faults sort of look the same.

Starting the debugger

As with `make` there are two different ways of invoking `gdb`. To start the debugger from the shell just type :

```
gdb <TargetName>
```

where `<TargetName>` is the name of the executable that you want to debug. If you do not specify a target then `gdb` will start without a target and you will need to specify one later before you can do anything useful.

As an alternative, from within `emacs` you can use the command `M-x gdb` which will then prompt you for the name of the target file. You cannot start an inferior `gdb` session from within `emacs` without specifying a target. The `emacs` window will then split between the `gdb` ‘window’ and a buffer containing the current source line.

Running the debugger

Once started, the debugger will load your application and its symbol table (which contains useful information about variable names, source code files, etc.). This symbol table is the map that the debugger reads as it is running your program.

^{vii} We recommend that you have fewer.

Warning: If you forget to specify the '-g' flag (debugging info.) when compiling your source files, this symbol table will be missing from your program and gdb (and you) will be "in the dark" as your program runs.

The debugger is an interactive program. Once started, it will prompt you for commands. The most common commands in the debugger are: setting breakpoints, single stepping, continuing after a breakpoint, and examining the values of variables.

Running the Program

<code>run</code>	Reset the program, run (or rerun) from the beginning. You can supply command-line arguments to 'run' the same way you can supply command-line arguments to your executable from the shell.
<code>step</code>	Run next line of source and return to debugger. If a subroutine call is encountered, follow into that subroutine.
<code>step count</code>	Run <i>count</i> lines of source.
<code>next</code>	Similar to <code>step</code> , but doesn't step into subroutines.
<code>finish</code>	Run until the current function/method returns.
<code>return</code>	Make selected stack frame return to its caller.
<code>jump address</code>	Continue program at specified line or address.

When a target application is first selected (usually on startup) the current source file is set to the file with the `main` function in it, and the current source line is the first executable line of the this function.

As you run your program, it will always be executing some line of code in some source file. When you pause the program (using a "breakpoint" or by hitting Control-C to interrupt), the "current target file" is the source code file in which the program was executing when you paused it. Likewise, the "current source line" is the line of code in which the program was executing when you paused it.

Breakpoints

You can use breakpoints to pause your program at a certain point. Each breakpoint is assigned an identifying number when you create it, and so that you can later refer to that breakpoint should you need to manipulate it.

A breakpoint is set by using the command `break` specifying the location of the code where you want the program to be stopped. This location can be specified in a variety of different ways, such as with the file name and either a line number or a function name within that file^{viii}. If the file name argument is not specified the file is assumed to be the current target file, and if no arguments are passed to `break` then the current source line will be the breakpoint. gdb provides the following commands to manipulate breakpoints:

<code>info break</code>	Prints a list of all breakpoints with numbers and status.
<code>break function</code>	Place a breakpoint at start of the specified function

^{viii} It is a good idea to specify lines that are really code, comments and whitespace will not do the right thing.

`break linenumber` Prints a breakpoint at line, relative to current source file.
`break filename:linenumber` Place a breakpoint at the specified line within the specified source file.

You can also specify an *if* clause to create a conditional breakpoint:

`break fn if expression` Stop at the breakpoint, only if *expression* evaluates to true. Expression is any valid C expression, evaluated within current stack frame when hitting the breakpoint.

`disable breaknum`
`enable breaknum` Disable/enable breakpoint identified by *breaknum*..

`delete breaknum` Delete the breakpoint identified by *breaknum*.

`commands breaknum` Specify commands to be executed when *breaknum* is reached. The commands can be any list of C statements or gdb commands. This can be useful to fix code on-the-fly in the debugger without re-compiling.

`cont` Continue a program that has been stopped.

For example, the commands:

```
break binky.c:120
break DoGoofyStuff
```

set a breakpoint on line 120 of the file `binky.c` and another on the first line of the function `DoGoofyStuff`. When control reaches these locations, the program will stop and give you a chance to look around in the debugger. If you're running `gdb` from `emacs`, it will load the source file and put an arrow (`->`) at the beginning of the line that is to be executed next.

Also, from inside `emacs`, you can set a breakpoint simply by go to the line of the file where you want to set the breakpoint, and hit `^x-[space]`.

`gdb` (and most other debuggers) provides mechanisms to determine the current state of the program and how it got there. The things that we are usually interested in are "where are we in the program?" and "what are the values of the variables around us?".

Examining the stack

To answer the question of "where are we in the program?", we use the `backtrace` command to examine the run-time stack. The run-time stack is like a "trail of breadcrumbs" in a program; each time a function call is made, a "crumb is dropped" (an RT stack frame is pushed). When a return from a function occurs, the corresponding RT stack frame is popped and discarded. These stack frames contain valuable information about where the function was called in the source code (line # and file name), what the parameters for the call were, etc.

`gdb` assigns numbers to stack frames counting from zero for the innermost (currently executing) frame. At any time `gdb` identifies one frame as the "selected" frame. Variable lookups are done with respect to the selected frame. When the program being debugged stops (at a breakpoint), `gdb` selects the innermost frame. The commands below can be used to select other frames by number or address.

`backtrace` Show stack frames, useful to find the calling sequence that produced a crash.

<code>frame <i>framenum</i></code>	Start examining the frame with <i>framenum</i> . This does not change the execution context, but allows to examine variables for a different frame.
<code>down</code>	Select and print stack frame called by this one.
<code>up</code>	Select and print stack frame that called this one.
<code>info args</code>	Show the argument variables of current stack frame.
<code>info locals</code>	Show the local variables of current stack frame.

Examining source files

Another way to find our current location in the program and other useful information is to examine the relevant source files. `gdb` provides the following commands:

<code>list <i>linenum</i></code>	Print ten lines centered around <i>linenum</i> in current source file.
<code>list <i>function</i></code>	Print ten lines centered around beginning of <i>function</i> (or <i>method</i>).
<code>list</code>	Print ten more lines.

The `list` command will show the source lines with the current source line centered in the range. (Using `gdb` from within `emacs` makes these command obsolete since it automatically loads the sources into `emacs` for you).

Examining data

It is also useful to answer the question, "what are the values of the variables around us?" In order to do so, we use the following commands to examine variables:

`print expression` Print value of *expression*. Expression is any valid C expression, can include function calls and arithmetic expressions, all evaluated within current stack frame.

`set variable = expression` Assign value of *variable* to *expression*. You can set any variable in the current scope. Variables which begin with `$` can be used as convenience variables in `gdb`.

`display expression` Print value of *expression* each time the program stops. This can be useful to watch the change in a variable as you step through code.

`undisplay` Cancels previous display requests.

In `gdb`, there are two different ways of displaying the value of a variable: a snapshot of the variable's current value and a persistent display for the entire life of the variable. The `print` command will print the current value of a variable, and the `display` command will make the debugger print the variable's value on every step for as long as the variable is 'live.' The desired variable is specified by using C syntax. For example:

```
print x.y[3]
```

will print the value of the fourth element of the array field named `y` of a structure variable named `x`. The variables that are accessible are those of the currently selected function's activation frame, plus all those whose scope is global or static to the current target file. Both the `print` and `display` functions can be used to evaluate arbitrarily complicated expressions, even those containing, function calls, but be warned that if a function has side-effects a variety of unpleasant and unexpected situations can arise.

Sometimes if you want to examine the contents of an array, you don't have to issue a `print` command for each element of the array. You can specify how many elements you want to print from a the `n`-th element in the array by using this extension to the `print` command:

```
print array[starting_point]@num_elements_to_print
```

This works on the `display` command also. For example, if you want to print ten elements of the array `x` starting from the 0th element, you would issue this:

```
print x[0]@10
```

Shortcuts

Finally, there are some things that make using `gdb` a bit simpler. All of the commands have short-cuts so that you don't have to type the whole command name every time you want to do something simple. A command short-cut is specified by typing just enough of the command name so that it unambiguously refers to a command, or for the special commands `break`, `delete`, `run`, `continue`, `step`, `next` and `print` you need only use the first letter.

Additionally, the last command you entered can be repeated by just hitting the <return key> again. This is really useful for single stepping for a range while watching variables change. If you're not running `gdb` from `emacs`, the up and down arrow keys will jump to the previous or next commands that you've issued, the left and right arrows will allow you to move through the command line for editing.

Miscellaneous

<code>editmode <i>mode</i></code>	Set <code>editmode</code> for <code>gdb</code> command line. Supported values for <i>mode</i> are <code>emacs</code> , <code>vi</code> , <code>dumb</code> .
<code>shell <i>command</i></code>	Execute the rest of the line as a shell command.
<code>history</code>	Print command history.

Debugging Strategy

If your program has been crashing spectacularly, you can just run the program by using the `run` command^{ix} right after you start the debugger. The debugger will catch the signal and allow you to examine the program (and hopefully find the cause and remedy).

More often the bug will be something more subtle. In these cases the "best" strategy is often to try to isolate the source of the bug, using breakpoints and checking the values of the program's variables before setting the program in motion using `run`, `step`, or `continue`. A common technique for isolating bugs is to set a breakpoint at some point before the offending code and slowly continuing toward the crash site examining the state of the program along the way.

^{ix} If your application takes command line arguments include these on the same line as the 'run' command.

Printing Your Source Files

There's a really neat way to print out hardcopies of your source files. Use a command called "enscript". Commonly, it's used at the UNIX command line as follows:

```
enscript -2GrPsweet3 binky.c lassie.c *.h
```

Where we want to print the two source files "binky.c" and "lassie.c", as well as all of the header files to printer sweet3. You can change these parameters to fit your needs.

Appendix A: UNIX^x/Emacs^{xi} Survival Guide

This handout summarizes many of the commands helpful for getting around on the Unix operating system and the Emacs editor.

Directory Commands

<code>cd <i>directory</i></code>	Change directory. If <i>directory</i> is not specified, goes to home directory.
<code>pwd</code>	Show current directory (<i>print working directory</i>)
<code>ls</code>	Show the contents of a directory. ls -a will also show files whose name begins with a dot. ls -l shows lots of miscellaneous info about each file
<code>rm <i>file</i></code>	Delete a file
<code>mv <i>old new</i></code>	Rename a file from old to new (also works for moving things between directories). If there was already a file named new, its previous contents are lost.
<code>cp <i>old new</i></code>	Creates a file named new containing the same thing as old. If there was already a file named new, its previous contents are lost.
<code>mkdir <i>name</i></code>	Create a directory
<code>rmdir <i>name</i></code>	Delete a directory. The directory must be empty.

Shorthand Notations & Wildcards

<code>.</code>	Current directory
<code>..</code>	Parent directory
<code>~</code>	Your home directory
<code>~<user></code>	Home directory of <i>user</i>
<code>*</code>	Any number of characters (not '!') Ex: *.c is all files ending in '.c'
<code>?</code>	Any single character (not '!')

Miscellaneous Commands

<code>cat <i>file</i></code>	Print the contents of <i>file</i> to standard output
<code>more <i>file</i></code>	Same as <code>cat</code> , but only a page at a time (useful for displaying)
<code>less <i>file</i></code>	Same as <code>more</code> , but with navigability (<code>less</code> is more)
<code>w</code>	Find out who is on the system and what they are doing
<code>ps</code>	List all your currently active processes
<code>jobs</code>	Show jobs that have been suspended

^x UNIX* is a registered trademark of AT&T

*UNIX** is a registered trademark of AT&T

**UNIX† is a registered trademark of AT&T

†...

^{xi} GNU EMACS is provided for free by the FREE SOFTWARE FOUNDATION, which write software and gives it away because they think that's how software should be.

<i>process</i> &	Runs a process in the background
%	Continue last job suspended (suspend a process with ^Z)
% <i>number</i>	Continue a particular job
kill <i>process-id</i>	Kill a process
kill -9 <i>process</i>	Kill a process with extreme prejudice
grep <i>exp files</i>	Search for an expression in a set of files
wc <i>file</i>	Count words, lines, and characters in a file
script	Start saving everything that happens in a file. type exit when done
lpr <i>file</i>	Print <i>file</i> to the default printer
lpr -Pinky <i>file</i>	Print <i>file</i> to the printer named <i>inky</i>
diff <i>file1 file2</i>	Show the differences between two files
telnet <i>hostname</i>	Log on to another machine
webster <i>word</i>	Looks up the given word in the dictionary. Works from most AIR machines, but is not standard UNIX (It actually talks to a NeXT...)

Getting Help

man <i>subject</i>	Read the manual entry on a particular subject
man -k <i>keyword</i>	Show all the manual listings for a particular keyword

History

history	Show the most recent commands executed
!!	Re-execute the last command
! <i>number</i>	Re-execute a particular command by number
! <i>string</i>	Re-execute the last command beginning with string
^ <i>wrong</i> ^ <i>right</i> ^	Re-execute the last command, substituting right for wrong
Ctrl-P	Scroll backwards through previous commands

Pipes

<i>a</i> > <i>b</i>	Redirect a's standard output to overwrite file b
<i>a</i> >> <i>b</i>	Redirect a's standard output to append to the file b
<i>a</i> >& <i>b</i>	Redirect a's error output to overwrite file b
<i>a</i> < <i>b</i>	Redirect a's standard input to read from the file b
<i>a</i> <i>b</i>	Redirect a's standard output to b's standard input

GNU EMACS

For the following "**^z**" means hit the "z" key while holding down the "ctrl" key. "**M-z**" means hit the "z" key while hitting the "META" or after hitting the "ESC" key.

Running Emacs

<code>emacs <filename></code>	Run emacs (on a particular file). Make sure you don't already have an emacs job running which you can just revive. Adding a ' & ' after the above command will run emacs in the background, freeing up your shell)
<code>^z</code>	Suspend emacs — revive with <code>%</code> command above
<code>^x^c</code>	Quit emacs
<code>^x^f</code>	Load a new file into emacs
<code>^x^v</code>	Load a new file into emacs and unload previous file
<code>^x^s</code>	Save the file
<code>^x-k</code>	Kill a buffer

Moving About

<code>^f</code>	Move forward one character
<code>^b</code>	Move backward one character
<code>^n</code>	Move to next line
<code>^p</code>	Move to previous line

(Note: the standard arrow keys also usually work.)

<code>^a</code>	Move to beginning of line
<code>^e</code>	Move to end of line
<code>^v</code>	Scroll down a page
<code>M-v</code>	Scroll up a page
<code>M-<</code>	Move to beginning of document
<code>^x-[</code>	Move to beginning of page
<code>M-></code>	Move to end of document
<code>^x-]</code>	Move to end of page
<code>^l</code>	Redraw screen centered at line under the cursor
<code>^x-o</code>	Move to other screen
<code>^x-b</code>	Switch to another buffer

Searching

<code>^s</code>	Search for a string
<code>^r</code>	Search for a string backwards from the cursor (quit both of these with <code>^f</code>)
<code>M-%</code>	Search-and-replace

Deletion

<code>^d</code>	Deletes letter under the cursor
-----------------	---------------------------------

`^k` Kill from the cursor all the way to the right
`^y` Yanks back all the last kills. Using the `^k ^y` combination you can get a cut-paste effect to move text around

Regions

Emacs defines a region as the space between the **mark** and the **point**. A mark is set with `^<spc>` (control-spacebar). The point is at the cursor position.

`M-w` Copy the region

`^w` Kill the region. Using `^y` will also yank back the last region killed or copied. This is what we used for "paste" back in the bad old mainframe days before there was "paste".

Screen Splitting

`^x-2` Split screen horizontally
`^x-3` Split screen vertically
`^x-1` Make active window the only screen
`^x-0` Make other window the only screen

Miscellaneous

`M-$` Check spelling of word at the cursor
`^g` In most contexts, cancel, stop, go back to normal command
`M-x goto-line <#>` Goes to the given line number
`^x-u` Undo
`M-x shell` Start a shell within emacs

Compiling

`M-x compile` Compile code in active window. Easiest if you have a makefile set up.
`^c ^c` Do this with the cursor in the compile window, scrolls to the next compiler error. YEAH!

Getting Help

`^h` Emacs help
`^h t` Run the emacs tutorial

Emacs does command completion for you. Typing `M-x <spc>` will give you a list of emacs commands. There is also a man page on emacs. Type `man emacs` in a shell.