

CSCI 315: Data Structures

Paul E. West, PhD

Department of Computer Science
Charleston Southern University

August 22, 2017

Linux File Structure


- In Linux, everything is a file. Well, almost!
- Programs can use disk files, serial ports, printers, and other devices in exactly the same way they would use a file.
- A file has a name and some properties, or “administrative information”
 - creation/modification date
 - its permissions
 - The properties are stored in the file’s inode
 - a special block of data in the file system
 - contains administrative information
 - contains the length of the file
 - where on the disk it’s stored

Directories

- Directories are files.
- A directory is a file that holds the inode numbers and names of other files.
- Each directory entry is a link to a file's inode; remove the filename and you remove the link.
- If the number of links to a file reaches zero, the inode and the data it references are no longer in use and are marked as free.
- This allows deletion when there are multiple links to the same file to be managed correctly.

Directoreis continued

- Files are arranged in directories, which may also contain subdirectories
- The / directory sits at the top of the hierarchy and contains all of the system's files in subdirectories
- The /home directory is a subdirectory of the root directory which is the home of all users
- /bin for system programs ("binaries")
- /etc for system configuration files,
- /lib for system libraries
- /dev for physical devices and provide the interface to those devices



Files and Devices

- Even hardware devices are very often represented (mapped) by files.
- You can mount a CD-ROM drive as a file:
- `# mount -t iso9660 /dev/hdc /mnt/cdrom`
- `# cd /mnt/cdrom`

/dev/console

- This device represents the system console
- Error messages and diagnostics are often sent to this device.
- On Linux, it's usually the “active” virtual console

/dev/tty

- The special file `/dev/tty` is an alias for the controlling terminal of a process
 - keyboard
 - screen
 - window
- `/dev/tty` allows a program to write directly to the user, without regard to which pseudo-terminal or hardware terminal the user is using

/dev/null

- This is the null device.
- All output written to this device is discarded.
- Unwanted output (aka a student's email complaint/rant) is often redirected to /dev/null.
- echo do not want to see this >/dev/null
- cp /dev/null empty_file

Introduction

- As a programmer or system administrator, you should know how to program under Linux
- We are going to learn
 - how to compile a program under Linux (gcc)
 - how to debug (gdb & ddd)
 - how to automate compilation (make)
 - how to perform memory analysis (valgrind)
 - how to create performance graphs (gnuplot - next set of slides.)

Compiling under C/C++

- We start with the simple case of all your source code in a single file.
- Try to generate a .c (NOT a .cpp) file as listed on the left hand side.

```
#include <stdio.h>

int main(int argc, char* argv[]) {
    printf("hello world\n");
    return 0;
}
```

Compile

- gcc test.c
- What file is generated?
- Name your compiled file by using
- gcc test.c -o test
- Run the generated executable file

Creating Debug Ready Code

- `cc -g test.c -o test`
- The '-g' flag tells the compiler to use debug info
- The compile file size is much larger
- We may still remove this debug information using the strip command
- `strip test`

Adding Optimizations

- The compiler can help improve the performance of your code via optimizations
- `cc -O test.c -o test`
- The '-O' flag tells the compiler to optimize the code.
- Usually can define an optimization level by adding a number to the '-O' flag

Getting Extra Compiler Warnings

- Error messages
 - Erroneous code that does not comply with the C standard
- Warnings
 - Codes that usually tend to cause errors during runtime
- Extra compiler warnings
 - useful to improve the quality of our source code
 - expose bugs that will really bug us later
- `cc -Wall test.c -o test`

Compiling a C++ program

```
#include <iostream>
```

```
int main(int argc, char* argv[]) {  
    std::cout << "hello world" << endl;  
    return 0;  
}
```


Compiling Multi Source Programs

- compile them
 - `cc main.c a.c b.c -o hello_world`
- Comments
 - external symbols need “extern” keyword
 - source file order becomes important
 - as program size increases so does compilation time

Limitation

- Even if we only make a change in one of the source files, all of them will be re-compiled when we run the compiler again.
- To overcome:
 - `cc -c main.cc`
 - `cc -c a.c`
 - `cc -c b.c`
 - `cc main.o a.o b.o -o hello_world`
- “-c” tells compiler only to create an object file, and not to generate a final executable file just yet
- The fourth command links the 3 object files into an executable file

Automating Program Compilation

- makefile is a collection of instructions that should be used to compile your program.
- Once you modify some source files, and type the command “make” (or “gmake” if using GNU’s make), your program will be recompiled using as few compilation commands as possible.

Makefile Structure

- Variable Definitions

- define values for variables for reuse

```
CFLAGS = -g -Wall  
SRCS = main.c file1.c file2.c  
CC = gcc
```

- Dependency Rules

- define under what conditions a given file needs to be re-compiled, and how to compile it.

```
main.o: main.c  
[tab]          gcc -g -Wall -c main.c
```

- if any of the files after : change,
 - then recompile
 - Note: You must use tabs in makefiles! (Spaces will NOT work.)
 - # is a comment

Single Source Makefile Example

```

# first you list your variable(s)
CC = gcc
# top-level rule to create the program.
# typically top rule is all (by convention)
all: main
# compiling the source file , main.o depends on main.c
main.o: main.c
    $(CC) -g -Wall -c main.c
# $(CC) uses value of CC variable , case sensitive
# linking the program , program name is main
main: main.o
    $(CC) -g main.o -o main
# cleaning everything that can be automatically recreated with "make".
# basically objects, the executable, and temp files
clean:
    rm -f main main.o

```

Multi Source file Example

```
# top-level rule to compile the whole program.
all: prog
# program is made of several source files.
prog: main.o file1.o file2.o
    gcc main.o file1.o file2.o -o prog
# rule for file "main.o".
main.o: main.c file1.h file2.h
    gcc -g -Wall -c main.c
# rule for file "file1.o".
file1.o: file1.c file1.h
    gcc -g -Wall -c file1.c
# rule for file "file2.o".
file2.o: file2.c file2.h
    gcc -g -Wall -c file2.c
# rule for cleaning files generated during compilations.
clean:
    rm -f prog main.o file1.o file2.o
```

Multi Source Make

- Commands can be anything, though usually they are gcc/g++ to compile or link
- Commands can be multiline, use tabs
- make param:
 - Runs the dependency for param
 - No param runs first label found (usually all)
 - make clean deletes files created by make

makedepend

- Program to generate dependencies
- Parameters are C/C++ files to scan
- It processes # directives to check includes
- Dependencies are critical, can do dependency graphs
- Can be added to Makefile:

```
depend: Makefile $(SRC)
    makedepend $(INCLUDES) $(SRC)
## end of Makefile
# DO NOT DELETE THIS LINE — make depend depends on it.
```

- Now make depend will run makedepend and add its results directly to Makefile. It doesn't generate the commands, but it does take care of finding dependencies for you.
- Many compilers automatically run dependency checks every time-never misses a dependency, but wasted work

configure

- Configure is part of the autoconf GNU system
- It makes Makefiles (ha!)
- Finds things (like say tcl) and updates Makefile variables with their locations
- Good for multi-platform Makefiles
- Not needed for this class

Debugging C/C++ Programs

- Before invoking the debugger, make sure you compiled your program with the "-g" flag (for gcc or g++)

```
gcc -g debug_me.c -o debug_me
```

```
gdb debug_me
```

```
[root@localhost root]# cc -g debug_me.c -o debug_me
[root@localhost root]# gdb debug_me
GNU gdb Red Hat Linux (6.0post-0.20040223.19rh)
Copyright 2004 Free Software Foundation, Inc.
GDB is free software, covered by the GNU General Public License, and you are
welcome to change it and/or distribute copies of it under certain conditions.
Type "show copying" to see the conditions.
There is absolutely no warranty for GDB. Type "show warranty" for details.
This GDB was configured as "i386-redhat-linux-gnu"...Using host libthread_db lib
rary "/lib/tls/libthread_db.so.1".

(gdb) █
```

Quitting

- q quits the debugger
- Only do this when you are done, not when you want to reload a program

Running the Program Inside a Debugger

- Once we invoked the debugger, we can run the program using the command "run".
- Can also just use the single letter r
- If the program requires command line parameters, we can supply them to the "run" command of gdb
run "hello, world" "goodbye, world"

Setting Breakpoints

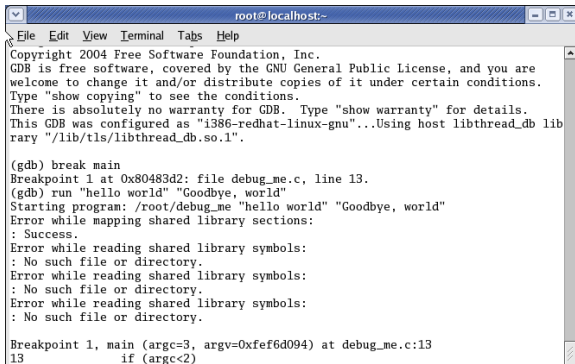
- A break point is a command for the debugger to stop the execution of the program before executing a specific source line.
- Specifying a specific line of code to stop in:
break 9 (assumes only 1 source file)
break debug_me.c:9
debug_me.c is the file name, 9 is the line number
- Specifying a function name, to break every time it is being called:
break main OR
b main
- Can also give a Boolean expression for conditional breakpoints

list

- list #
Shows that line number and next few lines
Hit return to show more lines
Can also just use letter l
- l debug_me:#

Stepping A Command At A Time

- Type command
break main
run "hello, world" "goodbye, world"



```
root@localhost:~  
File Edit View Terminal Tabs Help  
Copyright 2004 Free Software Foundation, Inc.  
GDB is free software, covered by the GNU General Public License, and you are  
welcome to change it and/or distribute copies of it under certain conditions.  
Type "show copying" to see the conditions.  
There is absolutely no warranty for GDB. Type "show warranty" for details.  
This GDB was configured as "i386-redhat-linux-gnu"...Using host libthread_db lib  
rary "/lib/tls/libthread_db.so.1".  
  
(gdb) break main  
Breakpoint 1 at 0x80483d2: file debug_me.c, line 13.  
(gdb) run "hello world" "Goodbye, world"  
Starting program: /root/debug_me "hello world" "Goodbye, world"  
Error while mapping shared library sections:  
: Success.  
Error while reading shared library symbols:  
: No such file or directory.  
Error while reading shared library symbols:  
: No such file or directory.  
Error while reading shared library symbols:  
: No such file or directory.  
  
Breakpoint 1, main (argc=3, argv=0xfef6d094) at debug_me.c:13  
13         if (argc<2)
```

Stepping A Command At A Time (continued...)

- Now we want to start running the program slowly, step by step.
- "next" (or n) - (step over) execute the current command, and stop again, showing the next command in the code to be executed.

```
Breakpoint 1, main (argc=3, argv=0xfef6d094) at debug_me.c:13
13         if (argc<2)
(gdb)
(gdb) next
18         for (argc--,argv++, i=1; argc>0; argc--, argv++,i++)
(gdb)
```


Stepping A Command At A Time (continued...)

- step (or s) - (step into functions) execute the current command, and if it is a function call - break at the beginning of that function
- step # - step # number of steps

```
Breakpoint 1, main (argc=3, argv=0xfef6d094) at debug_me.c:13
13         if (argc<2)
(gdb)
(gdb) next
18         for (argc--,argv++, i=1; argc>0; argc--, argv++,i++)
(gdb) step
20         print_string(i, argv[0]);
(gdb) █
```

Printing Variables And Expressions

- print (or p) the contents of a variable with a command like this:

- print i

```
(gdb) step
20                                     print_string(i, argv[0]);
(gdb) print i
$1 = 1
(gdb)
```

- You may also try to print more complex expressions, like "i*2", or "argv[3]", or "argv[argc]"

Examining The Function Call Stack

- Once we got into a break-point and examined some variables, we might also wish to see "where we are".
- This can be done using the "where" command

```
(gdb) where
#0  main (argc=2, argv=0xfef6d098) at debug_me.c:20
(gdb) █
```

Examining The Function Call Stack

- we can see contents of variables local to the calling function, or to any other function on the stack
frame 1
print i
- The "frame" command tells the debugger to switch to the given stack frame
- "0" is the frame of the currently executing function.

Debugging A Crashed Program

- Sometimes a program is will generate a core file containing its memory image when it crashes
- May also use backtrace (bt) shows series of functions calls to where we are
- Once we get such a core file, we can look at it by issuing the following command
gdb /path/to/program/debug_me core

Miscellaneous

- When you type run, it checks if executable has changed, but doesn't do recompilation
- kill terminates currently running code, but doesn't exit gdb
- help (or h) gives general help
- Help command gives help on some command

ddd

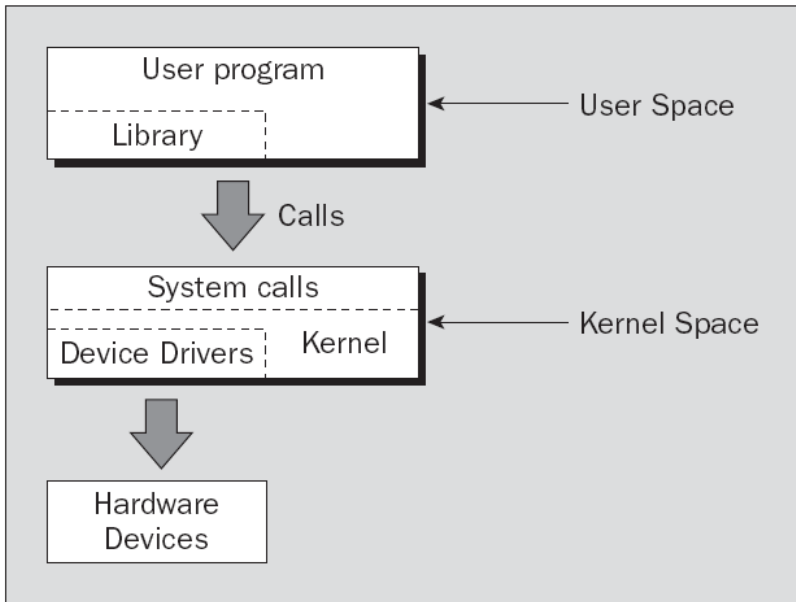
- GUI wrapper for gdb
- Also a wrapper for Java's jdb, Perl's and Python's pdb, among others
- apt-get install ddd

System Calls And Device Drivers

- We can access and control files and devices using system calls
- At the heart of the operating system, the kernel, are a number of device drivers.
- The low-level functions used to access the device drivers, the system calls, include:
 - open: Open a file or device
 - read: Read from an open file or device
 - write: Write to a file or device
 - close: Close the file or device
 - ioctl: Pass control information to a device driver

System Calls And Device Drivers

- The problem with using low-level system calls directly for input and output is that they can be very inefficient.
- Why?
 - Performance penalty in making a system call.
 - The hardware has limitations
- To provide a higher-level interface to devices and disk files, provides a number of standard libraries



What Even Is?

- A tool to perform
 - memory debugging
 - memory leak detection
 - memory profiling
- Valgrind accomplishes this by running your program inside of its virtual machine and capturing all your memory accesses/requests.

Memory debugging

● A successful run:

```
==2231== Memcheck, a memory error detector
==2231== Copyright (C) 2002–2013, and GNU GPL'd, by Julian Seward et al.
==2231== Using Valgrind-3.10.0 and LibVEX; rerun with -h for copyright info
==2231== Command: ./a.out
==2231==
==2231==
==2231== HEAP SUMMARY:
==2231==    in use at exit: 0 bytes in 0 blocks
==2231==    total heap usage: 0 allocs, 0 frees, 0 bytes allocated
==2231==
==2231== All heap blocks were freed — no leaks are possible
==2231==
==2231== For counts of detected and suppressed errors, rerun with: -v
==2231== ERROR SUMMARY: 0 errors from 0 contexts (suppressed: 0 from 0)
```

● Notice Valgrind keeps track of:

- (heap) memory used on exit
- How much heap memory was allocated & freed
- How many memory errors (out of bounds memory access) were detected.

● The “2231” is the process id, which for this class is unimportant

Memory Leak Detection

- Let's examine an obvious memory leak:

```
int main(int argc, char *argv[]) {  
    int *ints = new int[1024];  
  
    return 0;  
}
```

g++ test.c && valgrind ./a.out

```
==2476== HEAP SUMMARY:  
==2476==      in use at exit: 4,096 bytes in 1 blocks  
==2476==    total heap usage: 1 allocs, 0 frees, 4,096 bytes allocated  
==2476==  
==2476== LEAK SUMMARY:  
==2476==      definitely lost: 4,096 bytes in 1 blocks  
==2476==      indirectly lost: 0 bytes in 0 blocks  
==2476==      possibly lost: 0 bytes in 0 blocks  
==2476==      still reachable: 0 bytes in 0 blocks  
==2476==      suppressed: 0 bytes in 0 blocks
```

- Definitely lost means ... we lost ...
- Indirectly lost means, we lost and it could be hard to find.
- Possibly lost means Valgrind was not able to determine if the memory was deallocated or not.

Memory Usage detection

- A better version:

```
int main(int argc, char *argv[]) {  
    int *ints = new int[1024];  
  
    delete[] ints;  
    return 0;  
}
```

g++ test.c && valgrind ./a.out

```
==2591== HEAP SUMMARY:  
==2591==      in use at exit: 0 bytes in 0 blocks  
==2591==    total heap usage: 1 allocs, 1 frees, 4,096 bytes allocated  
==2591==  
==2591== All heap blocks were freed — no leaks are possible
```

- Notice Valgrind can tell us how much heap memory we are using even though we freed (deallocated) it.

Memory Access Error Detection

- Now lets use Valgrind to detect a off by one memory error.

```
int main(int argc, char *argv[]) {
    int *ints = new int[1024];

    for (int i = 0; i <= 1024; i++) {
        ints[i] = i;
    }

    delete[] ints;
    return 0;
}
```

Do you see the error?

```
==2615== Invalid write of size 4
==2615==    at 0x400673: main (in /home/pwest/t/a.out)
==2615==   Address 0x5a03040 is 0 bytes after a block of size 4,096 alloc'd
==2615==    at 0x4C298A0: operator new[](unsigned long) (vg_replace_malloc.c:389)
==2615==    by 0x40064E: main (in /home/pwest/t/a.out)
==2615==
==2615==
==2615== HEAP SUMMARY:
==2615==    in use at exit: 0 bytes in 0 blocks
==2615==   total heap usage: 1 allocs, 1 frees, 4,096 bytes allocate
```

- Well that tells us the error, but where is it?

Getting A Little More Help

- Compile with debug flags! (`g++ -g test.c && valgrind ./a.out`)

```
==2634== Invalid write of size 4
==2634==    at 0x400673: main (test.c:7)
==2634==   Address 0x5a03040 is 0 bytes after a block of size 4,096 alloc'd
==2634==    at 0x4C298A0: operator new[](unsigned long) (vg_replace_malloc.c:389)
==2634==   by 0x40064E: main (test.c:4)
==2634==
==2634==
==2634== HEAP SUMMARY:
==2634==    in use at exit: 0 bytes in 0 blocks
==2634==   total heap usage: 1 allocs, 1 frees, 4,096 bytes allocated
==2634==
==2634== All heap blocks were freed — no leaks are possible
```


Conclusion

- Valgrind can be used to detect memory leaks and memory access errors.
- Valgrind provides other tools to profile memory usage and cache usage, but those are beyond the scope of this class.
- Notice that Java has memory debugging built in, but C++ doesn't. Why?

Other Tools

- gprof: Used to profile
 - Profiling tells you where your program is spending the most time.
 - With that knowledge you can focus your energy on certain parts to speed up performance. (Why?)
- gnuplot: Generates simple graphs
 - gnuplot plugs in nicely with how we do things
 - You may use something else, if you **really** want too...
- Before gprof and gnuplot, we need to learn Bash!