

# Section 0: Tools

CS162

August 28-30, 2019

## Contents

<b>1</b>	<b>Make</b>	<b>2</b>
1.1	More details about Make . . . . .	2
<b>2</b>	<b>Git</b>	<b>3</b>
2.1	Commands to know . . . . .	3
<b>3</b>	<b>GDB: The GNU Debugger</b>	<b>4</b>
3.1	Commands to know . . . . .	4
<b>4</b>	<b>Debugging programs on Linux</b>	<b>5</b>
4.1	strace and ltrace . . . . .	5
4.2	catchsegv . . . . .	5
4.3	top . . . . .	6
4.4	ps . . . . .	6
4.5	pidof and pgrep . . . . .	6
4.6	proc . . . . .	6
4.7	lsof . . . . .	7
4.8	netstat and ip . . . . .	7
4.9	kill, pkill, and killall . . . . .	7
<b>5</b>	<b>Debugging Example</b>	<b>8</b>

Tools are important for every programmer. If you spend time learning to use your tools, you will save even more time when you are writing and debugging code. This section will introduce the most important tools for this course.

## 1 Make

GNU Make is program that is commonly used to build other programs. When you run `make`, GNU Make looks in your current directory for a file named `Makefile` and executes the commands inside, according to the makefile language.

```
my_first_makefile_rule:
    echo "Hello world"
```

The building block of GNU Make is a **rule**. We just created a rule, whose **target** is `my_first_makefile_rule` and **recipe** is `echo "Hello world"`. When we run `make my_first_makefile_rule`, GNU Make will execute the steps in the recipe and print “Hello world”.

Rules can also contain a list of **dependencies**, which are other targets that must be executed before the rule. In this example, the `task_two` rule has a single dependency: `task_one`. If we run “`make task_two`”, then GNU Make will run `task_one` and then `task_two`.

```
task_one:
    echo 1
task_two: task_one
    echo 2
```

### 1.1 More details about Make

- If you just run `make` with no specified target, then GNU Make will build the first target.
- By convention, target names are also file names. If a rule’s file exists and the file is **newer** than all of its dependencies, then GNU Make will skip the recipe. If a rule’s file does not exist, then the timestamp of the target would be “the beginning of time”. Otherwise, the timestamp of the target is the **Modification Time** of the corresponding file.
- When you run “`make clean`”, the “clean” recipe is executed every time, because a corresponding file named “clean” is never actually created. (You can also use the `.PHONY` feature of the makefile language to make this more robust.)
- Makefile recipes must be indented with tabs, not spaces.
- You can run recipes in parallel with “`make -j 4`” (specify the number of parallel tasks).
- GNU Make creates automatic rules if you don’t specify them. For example, if you create a file named `my_program.c`, GNU Make will know how to compile it if you run “`make my_program`”.
- There are many features of the makefile language. Special variables like `$$` and `$$<` are commonly used in Makefiles. Look up the documentation online for more!

Pintos, the educational operating system that you will use for projects, has a complex build system written with Makefiles. Understanding GNU Make will help you navigate the Pintos build system.

## 2 Git

Git is a distributed revision control and source code management (SCM) system with an emphasis on speed, data integrity, and support for distributed, non-linear workflows. GitHub is a Git repository hosting service, which offers all of the distributed revision control and SCM functionality of Git as well as adding many useful and unique features.

In this course, we will use Git and GitHub to manage all of our source code. It's important that you learn Git, but NOT just by reading about it. If you're unfamiliar with Git, go to <https://try.github.io/> and try it out yourself.

### 2.1 Commands to know

- **git init**  
Create a repository in the current directory
- **git clone <url>**  
Clone a repository into a new directory
- **git status**  
Show the working tree status
- **git pull <repo> <branch>**  
Fetch from and integrate with another repository or a local branch
- **git push <repo> <branch>**  
Update remote refs along with associated objects
- **git add <file(s)>**  
Add file contents to the index
- **git commit -m <commit message>**  
Record changes to the repository with the provided commit message
- **git branch**  
List or delete branches
- **git checkout**  
Checkout a branch or paths to the working tree
- **git merge**  
Join two or more development histories together
- **git rebase**  
Reapply commits on top of another base
- **git diff [--staged]**  
Show a line-by-line comparison between the current directory and the index (or between the index and HEAD, if you specify --staged).
- **git show [--format=raw] <tree-ish>**  
Show the details of anything (a commit, a branch, a tag).
- **git reset [--hard] <tree-ish>**  
Reset the current state of the repository
- **git log**  
Show commits on the current branch

- **git relog**  
Show recent changes to the local repository

When you begin working on group projects for this course, an advanced knowledge of Git will help you manage your team's code. Git is especially useful when multiple developers want to work on the same codebase simultaneously.

## 3 GDB: The GNU Debugger

GDB is a debugger that supports C, C++, and other languages. You will not be able to debug your projects effectively without advanced knowledge of GDB, so make sure to familiarize yourself with GDB as soon as possible.

### 3.1 Commands to know

- **run, r**  
Start program execution from the beginning of the program. Also allows argument passing and basic I/O redirection.
- **quit, q**  
Exit GDB
- **kill**  
Stop program execution.
- **break, break x if condition, clear**  
Suspend program at specified function (e.g. “`break strcpy`”) or line number (e.g. “`break file.c:80`”). The “clear” command will remove the current breakpoint.
- **step, s**  
If the current line of code contains a function call, GDB will step into the body of the called function. Otherwise, GDB will execute the current line of code and stop at the next line.
- **next, n**  
Execute the current line of code and stop at the next line.
- **continue, c**  
Continue execution (until the next breakpoint).
- **finish**  
Continue to end of the current function.
- **print, p**  
Print value stored in variable.
- **call**  
Execute arbitrary code and print the result.
- **watch, rwatch, awatch**  
Suspend program when condition is met. i.e. `x > 5`.
- **backtrace, bt, bt full**  
Show stack trace of the current state of the program.

- **disassemble**  
Show an assembly language representation of the current function.
- **set follow-fork-mode <mode>** (Mac OS does not support this)  
GDB can only debug 1 process at a time. When a process forks itself (creates a clone of itself), we can follow the parent (original) or the child (clone). The mode argument can be either **parent** or **child**.

The **print** and **call** commands can be used to execute arbitrary lines of code while your program is running! You can assign values or call functions. For example, “**call close(0)**” or “**print i = 4**”. (You can actually use **print** and **call** interchangeably most of the time.) This is one of the most powerful features of gdb.

## 4 Debugging programs on Linux

Debuggers and print statements are not the only ways to debug programs. Your student VM comes with a bunch of tools that provide information about running programs. Here are some of them that you will find useful for completing homework in this course.

### 4.1 strace and ltrace

All programs use standard library functions like **printf**, and most library functions involve **system calls** that interact with the kernel. This is true whether your program is written in C, Python, Java, or any other language. **strace** and **ltrace** are two programs that print out system calls (**strace**) and standard library function calls (**ltrace**) used by a program.

```
# Run 'ls files/', but print out system calls
strace ls files/

# Run 'ls files/', but print out library calls
ltrace ls files/

# Print out system calls of a running process with process ID = 1234
strace -p 1234

# Print out library calls of a running process with process ID = 1234
ltrace -p 1234
```

You can use **strace** and **ltrace** to debug a hung process, debug performance issues, or peek into the internals of somebody else’s program if you’re curious. You can add the **-f** flag to make **strace** and **ltrace** follow the descendants of the original process.

### 4.2 catchsegv

**catchsegv** will give you helpful information when your program encounters a segmentation fault.

```
catchsegv ./my_buggy_program
```

There are 3 parts to **catchserv** output. The 1st part prints out the values of all the CPU registers when the program crashed. The 2nd part prints out a stack trace of the crash, but only if you compiled with debug flags (**-g**). The 3rd part prints out the memory map of the crash—we will discuss shared memory in this course. **catchsegv** also handles aborts, traps, and other abnormal program crashes.

### 4.3 top

`top` is a text user interface for viewing the most active processes on your system. Press `q` to quit. Press `c` to show full commands. Press `<` or `>` to change the sorted column. Here are some of the most useful columns of `top`:

- PID – The process ID
- PR, NR – The priority value and nice value
- VIRT – The amount of virtual memory used by the process (KiB)
- RES – The amount of resident memory (actually located in RAM) used by the process (KiB)
- SHR – The amount of memory that is sharable (shared libraries, mapped files, etc) (KiB)
- S – The state of the process (S for sleeping, R for running, T for stopped, Z for zombie)
- %CPU, %MEM – The percentage of CPU and RAM used by the process
- TIME+ – The amount of “CPU time” used by the process. The unit for this is seconds. If a process is idle, it is using 0 CPU time.
- COMMAND – The name of the process

### 4.4 ps

`ps` prints out a list of running processes. By default, `ps` only prints the processes that are part of your own console session, which probably only includes `bash` and `ps` itself. You can use “`ps -eLf`” or “`ps auxww`” to get way more detail. (There’s a long history about why one of these commands requires a dash.)

The `ps` command contains mostly the same information as `top`, but it is easier to filter and search, because it only appears once and shows all the processes.

The **PPID** column of “`ps -eLf`” tells you the **parent process ID**, which is the process ID of a process’s parent. The **NLWP** column shows the number of light-weight processes (also known as threads) that belong to a particular process.

### 4.5 pidof and pgrep

`pidof <process name>` gets you the process ID of a single running process. For example, `pidof vim` will get me the process ID of my vim process. If there are multiple matches, then all the process IDs are printed. It’s useful as an argument to another program. For example, I could attach `gdb` to my `httpserver` process by running “`gdb -p $(pidof httpserver)`”.

You can also use `pgrep <process name>`, which doesn’t require the full process name, but only a part of it. For example, “`pgrep a`” would give me the PID of all processes that contain ‘a’ in their name.

### 4.6 proc

The `/proc` directory on Linux is a gold mine of information. You can run `ls /proc/1` to get information about the process with `PID = 1`. Many of the tools mentioned earlier are implemented by reading the `/proc` directory, so if you need more detail about a process, look there! The `/proc/self` directory is a shortcut to the current process. (If you run `ls /proc/self`, then you will get information about the `ls` process!)

Some useful parts of `proc`:

- `cmdline` – The command used to start the process
- `cwd` – The current working directory of the process
- `environ` – The environment variables of the process
- `exe` – The program that is running in the process
- `fd/` – A list of the file descriptors of the process
- `fdinfo/` – More information about the file descriptors
- `maps` – The memory map of the process
- `net` – Details about network configuration
- `status` – Status of the process

## 4.7 `lsof`

`lsof` gets you information about the files, processes, and network connections on your system. Each line of output from `lsof` corresponds to an “open file”, which can correspond to a process’s current directory, an actual file that is in use by a process, a shared library, a network connection, or a chunk of shared memory. Here are a few useful ways to use it:

- `lsof /lib/x86_64-linux-gnu/libc.so.6` – List instances of programs using the C standard library
- `lsof -p 1234` – List open files for the process with process ID = 1234
- `lsof -i tcp:80` – List uses of TCP port 80 (http)

## 4.8 `netstat` and `ip`

`netstat` and `ip` are two commands that help you debug the network connections on your system. Here are a few useful ways to use them:

- You can run “`ip addr`” or “`ifconfig`” to get the IP addresses and network interfaces of the system.
- You can run “`ip route`” or “`netstat -nr`” to get the routing table.
- You can run “`netstat -nt`” to get the list of current TCP connections on the system.
- You can run “`netstat -ntl`” to get the TCP ports that are listening for new connections.

Several homework assignments in this course will require the use of network connections in your programs. We will introduce more tools for testing network connections when those homework assignments are released.

## 4.9 `kill`, `pkill`, and `killall`

To kill a process, you can use “`kill <PID>`”. You can also attach a signal to the command, like “`kill -TERM <PID>`” or “`kill -QUIT <PID>`”.

`pkill` will kill all processes that match a pattern. For example, “`pkill v`” will kill all processes with ‘v’ in the name. `killall` does the same thing as `pkill`, but requires an exact match.

## 5 Debugging Example

Take a moment to read through the code for `asuna.c`. It takes in 0 or 1 arguments. If an argument is provided, `asuna` uses quicksort to sort all the chars in the argument. If no argument is provided, then `asuna` uses a default string to sort.

```

1 int partition(char* a, int l, int r){
2     int pivot, i, j, t;
3     pivot = a[l];
4     i = l; j = r+1;
5
6     while(1){
7         do
8             ++i;
9         while( a[i] <= pivot && i <= r );
10        do
11            --j;
12        while( a[j] > pivot );
13        if( i >= j )
14            break;
15        t = a[i];
16        a[i] = a[j];
17        a[j] = t;
18    }
19    t = a[l];
20    a[l] = a[j];
21    a[j] = t;
22    return j;
23 }

1 void sort(char a[], int l, int r){
2     int j;
3
4     if(l < r){
5         j = partition(a, l, r);
6         sort(a, l, j-1);
7         sort(a, j+1, r);
8     }
9 }

1 void main(int argc, char** argv){
2     char* a = NULL;
3     if(argc > 1)
4         a = argv[1];
5     else
6         a = "Asuna is the best char!";
7     printf("Unsorted: \"%s\"\n", a);
8     sort(a, 0, strlen(a) - 1);
9     printf("Sorted:   \"%s\"\n", a);
10 }

```

When `asuna` is run, we get the following output:

```

$ ./asuna "Kirito is the best char!"
Unsorted: "Kirito is the best char!"
Sorted   : "   !Kabceehhiiiorrssttt"

```

```

$ ./asuna
Unsorted: "Asuna is the best char!"
Segmentation fault (core dumped)

```

Use the debugging tools to find why `asuna.c` crashes when no arguments are provided.

*Hint: `catchsegv` is a great place to start for debugging seg faults.*



The first step in debugging a seg fault is often times seeing if `catchsegv` can give you a useful backtrace. To use `catchsegv`, make sure to compile `asuna` with

```
$ gcc -g -no-pie asuna.c -o asuna
```

Turning off position independent executable will allow `catchsegv` to provide a more detailed backtrace. Then run `catchsegv` with

```
$ catchsegv ./asuna
... register dump omitted ...
Backtrace:
/home/vagrant/code/personal/disc/section0/asuna.c:20(partition) [0x4005e5]
/home/vagrant/code/personal/disc/section0/asuna.c:34(sort) [0x400679]
/home/vagrant/code/personal/disc/section0/asuna.c:48(main) [0x40071d]
/lib/x86_64-linux-gnu/libc.so.6(__libc_start_main+0xe7) [0x7fccf82ecb97]
???(?_start) [0x40047a]
... memory map omitted ...
```

You can also get the same backtrace with `gdb` using the `backtrace` or `bt` command immediately after the segfault occurs. Notice the backtrace points to an error in the `partition` function, specifically the line `a[i] = a[j]`. We can inspect this bug closer now that we know where its located by using `gdb` or `cgdb`. We can either set the breakpoint to be on `partition` or the actual faulting line.

```
$ gdb ./asuna
(gdb) b asuna.c:20 # set a breakpoint on the faulting line
(gdb) r # runs the program until the breakpoint
(gdb) n # runs the next line, which segfaults
```

At this point, notice that

1. This line performs 2 operations: a read from `a[j]` and a write to `a[i]`.
2. Earlier in the program we already execute a `a[j]` in `partition:12`.
3. If we run `asuna` with the default argument passed in, no segfault occurs.

The fact that #1 and #2 are simultaneously true points to a problem with the write to `a[i]`, which is most likely a memory issue. #3 implies that memory is somehow different when using a default argument vs a user provided argument. In `gdb`, we can print the address of the string `a` when using a default arg compare a user provided arg.

```
(gdb) print a
$1 = 0x4007c4 "Asuna is the best char!"
(gdb) r "Test user argument" # rerun the program with a user arg
The program being debugged has been started already.
Start it from the beginning? (y or n) y
(gdb) print a
$2 = 0x7fffffff6fa "Test user argument"
```

Notice how the address of the default arg is so much lower than that of the user provided arg. This is because the default argument is in the static region of the program. The segfault occurs because memory in the static region cannot be modified. When a string is declared as part of the program such as in `main:6`, that string is compiled into the code and stored in static. See this [Stackoverflow](#) for a more detailed explanation of this bug.

Below we provide a cleaned up and fixed version of the the same program. Our solution is to malloc an array on the heap for the argument to `partition` and `strcpy` the string into that array.

```

1 void swap (char* arr, int first, int second) {
2     char temp = arr[first];
3     arr[first] = arr[second];
4     arr[second] = temp;
5 }
6
7 int partition(char* arr, int left_bound, int right_bound){
8     int pivot = arr[left_bound];
9     // Initialize to starting bounds we won't use
10    int left_loc = left_bound;
11    int right_loc = right_bound + 1;
12
13    while(left_loc < right_loc){
14        // Make forward progress on every iteration
15        // so use do while loops
16        do {
17            left_loc++;
18            // Find the leftmost elem greater than the pivot
19        } while (left_loc <= right_bound && arr[left_loc] <= pivot);
20
21        // Make forward progress on every iteration
22        // so use do while loops
23        do {
24            right_loc--;
25            // Find the rightmost elem less than the pivot
26        } while (right_loc > left_loc && arr[right_loc] > pivot);
27        // If there are elements to switch swap them
28        if(left_loc < right_loc) {
29            swap (arr, left_loc, right_loc);
30        }
31    }
32    swap (arr, left_bound, right_loc);
33    return right_loc;
34 }
35
36 void sort(char* arr, int left_bound, int right_bound){
37     if(left_bound < right_bound){
38         // divide and conquer
39         int split_point = partition(arr, left_bound, right_bound);
40         sort(arr, left_bound, split_point-1);
41         sort(arr, split_point+1, right_bound);
42     }
43 }
44
45 void main(int argc, char** argv){
46     const char* no_args = "Asuna is the best char!";
47     char* arr = NULL;
48     if(argc > 1) {
49         arr = malloc (strlen (argv[1]) * sizeof (char));
50         strcpy (arr, argv[1]);
51     } else {
52         arr = malloc (strlen (no_args) * sizeof (char));
53         strcpy (arr, no_args);

```

```
54     }
55     printf("Unsorted: \"%s\"\n", arr);
56     sort(arr, 0, strlen(arr) - 1);
57     printf("Sorted:   \"%s\"\n", arr);
58     // Really not necessary because this is main but
59     // might as well free all your mallocs
60     free (arr);
61 }
```