

Basic Unix

2.1 What is Unix?

Unix is an *operating system*, by which we broadly mean the suite of software that make the computer work. Specifically, the operating system provides *services* for the program they run. Unix (and its cousin Linux) is used by the terminals and multi-user servers in most physics/astronomy departments. For people with Windows on their laptop, we will use Cygwin to provide Unix functionality.

The Unix operating system is made up of three parts: the kernel; the shell; and the programs.

2.1.1 The Kernel

The *kernel* of Unix is the hub of the operating system: it allocates time and memory to programs and handles the file store and communications in response to system calls (requests from programs for the kernel to do things on behalf of the program, e.g. print characters on the screen, write data to disk, etc).

2.1.2 The Shell

The *shell* acts as an interface between the user and the kernel. When a user logs in, the login program checks the username and password, and then starts another program called the *shell*. The shell is a command line interpreter (CLI). It interprets the commands the user types in and arranges for them to be carried out. The commands themselves are usually programs: when they terminate, the shell gives the user another prompt to indicate it is ready to accept more commands.

There are many different shells, one of the earliest being the Bourne shell (`/bin/sh`) which is pretty much available on all Unix distributions. We will be showing you Bash (Bourne Again SHell¹) (`/bin/bash`). Bash is an sh-compatible shell that incorporates useful features from other more modern shells.

2.1.3 Files and Processes

Everything in Unix is either a file or a process.

A process is an executing program identified by a unique PID (process identifier).

A file is a collection of data. Files can be created by users using text editors, running programs, or copying existing files from elsewhere on the system, or across the network etc. Notice that the idea of a file does not talk specifically about being on disk, it could be in memory, on disk, being transmitted across the network or being generated on the fly by another program or the operating system.

2.2 Working with Files and Directories

Just like in the Windows environment, files are stored in a hierarchical *directory* structure. The location of a particular file (or *subdirectory*) is described using a *path*. Under Windows, the hierarchy starts from a drive letter (for example `C:\`) but in Unix the hierarchy starts from a single slash (`/`) called the *root* directory. An *absolute path* is one that starts from the top of the hierarchy (i.e. begins with a slash). Alternatively, a *relative path* starts from the current directory.

An absolute path points to the same place regardless of what the current directory is. A relative path is relative to your current location.

Don't directories break the *everything is a file or process* rule? Directories are in fact special files which link

¹Welcome to bad Unix naming humour!

filenames and other information about a file, with the file data itself. It is possible (although not advisable) to open a directory just like any other file on Unix systems.

2.2.1 pwd

Each process (a running program) has a *working directory* (also called the *current working directory*) from which all relative paths are interpreted.

The **pwd** command prints the current path and thus enables you to work out where you are in the filesystem. For example, if you run **pwd** when you first log in, you will find you are in your *home directory*:

```
1 % pwd
2 /Users/tara
3 %
```

Of course your home directory path will be different, other examples are:

```
1 /cygdrive/c/Users and Documents/tara
2 /home-astrop/tara
3 /usr/extern/tara
```

2.2.2 ls

When you first login, your current working directory is your home directory. Your home directory has the same name as your user name, for example, **tara**, and it is where your personal files and subdirectories are saved.

To find out what is in your home directory, type:

```
1 % ls
2 README
3 %
```

The **ls** command lists the contents of your current working directory.

If the home directory is empty no files will be shown (and you will just get a prompt back on the next line). Usually there are some standard files inserted by the System Administrator when your account was created.

ls does not, in fact, cause *all* the files in your home directory to be listed, but only those ones whose name does not begin with a dot (**.**) Files beginning with a dot (**.**) are known as *hidden files* and usually contain important program configuration information. They are hidden because you should not change them unless you are very familiar with Unix!

To list all files in your home directory including those whose names begin with a dot, type:

```
1 % ls -a
2 .
3 ..
4 .login
5 README
6 %
```

ls is an example of a command which can take options: **-a** is an example of an option. The options change the behaviour of the command. There are online manual pages that tell you which options a particular command can take, and how each option modifies the behaviour of the command. (See later in these lab notes)

2.2.3 mkdir

We will now make a subdirectory in your home directory to hold the files you will be creating and using in this lab. To make a subdirectory called **unixstuff** in your current working directory type:

```
1 % mkdir unixstuff
2 %
```

To see the directory you have just created, type:

```
1 % ls
2 README unixstuff
```

```
3 %
```

2.2.4 cd

The *change directory* command `cd <directory>`, changes the current working directory to `<directory>`. The current working directory may be thought of as the directory you are “in”, i.e. your current position in the file-system tree.

To change to the directory you have just made, type:

```
1 % cd unixstuff
2 %
```

Type `ls` to see the contents (which should be empty) and `pwd` to see the absolute path.

```
1 % ls
2 % pwd
3 /Users/tara/unixstuff
4 %
```

2.2.5 The Special directories . and ..

Using `ls -a` will always display the two directories `.` and `..`. They represent the *current directory* and its *parent directory* respectively. Clearly, these special directories are relative paths because they depend (by definition) on the current working directory.

These can be used with `cd` to navigate the directories:

```
1 % pwd
2 /Users/tara/unixstuff
3 % cd .
4 % pwd
5 /Users/tara/unixstuff
6 % cd ..
7 % pwd
8 /Users/tara
9 %
```

As you can see, `cd ..` takes you up to the parent directory.

When we type `cd .` we do not move as we are asking to move to the directory we are already in. It might not seem very useful yet, but there are lots of places later on where we will need to refer to the current directory.

2.2.6 More about directories and pathnames

```
1 % pwd
2 /Users/tara
3 % mkdir foo/bar
4 mkdir: cannot create directory 'foo/bar': No such file or directory
5 % mkdir unixstuff/backups
6 % ls unixstuff
7 backups
8 % mkdir -p foo/bar
9 % ls foo
10 bar
11 %
```

We cannot create nested directories (directories within other directories) without first creating the necessary parent directories. The `-p` option can be used to force the building of parent directories first. Now we have two directories within our home directory (`unixstuff` and `foo`). The `unixstuff` directory has a `backups` subdirectory, and `foo` has a `bar` subdirectory.

```
1 % ls backups
2 ls: backups: no such file or directory
3 %
```

We can't see inside the **backups** directory directly from our home directory, because it is only accessible from inside the **unixstuff** directory. However, we explicitly put parent directories using slashes (a relative path) like:

```

1 % cd unixstuff/backups
2 % pwd
3 /Users/tara/unixstuff/backups
4 % ls ..
5 backups
6 % ls ../../..
7 README      foo          unixstuff
8 % ls ../../foo
9 bar
10 %

```

Notice we can also use the `..` directory with the slashes to get access to the grandparent directory, and then other directories from there.

`~` (the *tilde* or *twiddle* character²) is another special directory that represents your home directory. To get to another user's home directory (if they have made it accessible to others), you would use `~<login>`.

```

1 % ls ~
2 unixstuff
3 % cd ~/unixstuff
4 % pwd
5 /Users/tara/unixstuff
6 %

```

Notice that we can combine `~` with other directories using the slash.

2.3 Copy, Moving and Deleting Files and Directories

2.3.1 cp

This is used to copy a file or directory to a new location:

```

1 % cd ~
2 % touch afile
3 % ls
4 README      afile      unixstuff
5 % cp afile unixstuff
6 % ls unixstuff
7 afile      backups
8 %

```

touch is used to create an empty file (containing no characters), but you could create a file like this just as easily with a text editor (e.g. `pico`, `vi`, `emacs`, `nedit`).

Here we have copied the file **afile** into the **unixstuff** directory. So now a copy of **afile** exists in your home directory and also in the **unixstuff** directory.

The file can also be copied to different filenames. Here we make a copy of **afile** called **anotherfile**:

```

1 % cp afile anotherfile
2 % ls
3 README      afile      anotherfile foo          unixstuff
4 %

```

As well as specifying a named directory, you can also copy files to the current directory using dot:

```

1 % cd unixstuff
2 % cp ../anotherfile .
3 % ls
4 afile      anotherfile backups
5 %

```

²Tilde can be a bit hard to find on some keyboards – typically it is on the top left above the tab key.

If we want to copy directories and all of their subdirectories *recursively*, the `-r` flag should be used to copy recursively:

```
1 % cd
2 % cp -r unixstuff moreunixstuff
3 % ls moreunixstuff
4 afile          anotherfile  backups
5 %
```

Notice that the `cd` command without a directory argument returns you to your home directory (just like `cd ~`). We have then created a complete copy of `unixstuff` and its subdirectory `backups`.

2.3.2 mv

This is used to move (or rename) files and directories. Note, unlike MSDOS and the Windows command shell (`cmd.exe`) there is no separate rename command:

```
1 % ls
2 README          anotherfile  moreunixstuff
3 afile           foo          unixstuff
4 % mv afile foo
5 % mv anotherfile yetanotherfile
6 % mv moreunixstuff lessunixstuff
7 % ls
8 README          foo          lessunixstuff  unixstuff     yetanotherfile
9 % ls foo
10 afile bar
11 %
```

Here we have renamed the file `anotherfile` to `yetanotherfile` and `moreunixstuff` to `lessunixstuff`. We have also moved `afile` from our current directory into the `foo` directory.

While we are here, notice that the `ls` command tries to format the filenames in the most readable manner. This will change depending on the size of the shell window, and the length and number of filenames in the directory being listed.

2.3.3 rm

This is used to delete files.

```
1 % ls
2 README          foo          lessunixstuff  unixstuff     yetanotherfile
3 % rm yetanotherfile
4 % ls
5 README          foo          lessunixstuff  unixstuff
6 %
```

Here we have deleted the file `yetanotherfile`. NB: unlike Windows there is no trash can under Unix, once a file is deleted it is gone forever. The only way of getting back a file is to have a backup copy of it somewhere else on the system. Most (all?) departments will backup user accounts every evening, so it is possible to retrieve your previous days work by speaking to your System Administrator.

```
1 % rm lessunixstuff
2 rm: lessunixstuff is a directory
3 % rm -r lessunixstuff
4 % ls
5 README          foo          unixstuff
6 %
```

We cannot use this to delete directories unless we use the `-r` flag. Be careful using this as it will delete everything in that directory, i.e. all of the subdirectories and the files they contain. In fact, be **very careful with `rm -r`**!

2.3.4 rmdir

This is used to delete directories. It will only work when they are empty.

```

1 % rmdir foo
2 rmdir: directory "foo": Directory not empty
3 % rm foo/afile
4 % rmdir foo/bar
5 % ls foo
6 % rmdir foo
7 % ls
8 README      unixstuff
9 %

```

This series of commands deleted the file `afile` and the directory `bar` from within the `foo` directory, so we could then delete the `foo` directory using `rmdir`. You can see why using `rm -r` is tempting, but be careful!

2.4 Displaying Files

For these exercises you will need to download the word list from the BasicUnix wiki page. This file contains a list of standard English words. It is often found in `/usr/dict/words` or `/usr/share/dict/words` on modern Linux systems. Once you have downloaded the file, put it in your `unixstuff` directory.

2.4.1 clear

This command will clear the screen and is sometimes useful when viewing files or listing lots of directories. It doesn't clear the entire history, only the current screenful. Commands prior to that will still be visible if you scroll the terminal window up.

2.4.2 cat

This will print the given file(s) to the standard output (screen).

```

1 % cd ~/unixstuff
2 % cat words.txt
3 10th
4 1st
5 ...
6 Zurich
7 zygote
8 %

```

You should have seen a lot of words fly by on the screen with this command — all the words contained in this word list. If multiple files are given to `cat` they will be printed one after the other in order.

2.4.3 less

`less` is a more advanced text viewer. It allows the navigation a file or data read from standard input (normally the keyboard). It is a more advanced version of another text viewer, called `more`³, which didn't have the capability of scrolling backwards through a file.

Here are some of the supported commands:

arrow keys or Enter	up or down one line at a time
PageUp/Down	up or down one page at a time
b or SpaceBar	
g	beginning of the file
G	end of the file
/<pattern>	find a given pattern <pattern>
n	find the next matching instance of the pattern
q	quit out of less
h	get help on the available commands

`less` can be used as a direct substitute for `cat`:

```

1 % less words.txt
2 10th

```

³see what we mean about Unix humour!

```

3 1st
4 2nd
5 ... [to end of page]
6 %

```

But it can also be used in the following way:

```

1 % cat words.txt | less
2 10th
3 1st
4 2nd
5 ... [to end of page]
6 %

```

Both these have the same effect. The second idiom can be used to view the output of any program that produces multiple screenfuls of output. Here the vertical bar (|) is a *pipe*, indicating that the output from the first program should not be output to the screen, but rather should be used as input to the second program (instead of the keyboard). This means `cat` is providing input to `less` rather than `less` directly reading it from the file. We will see why this is useful shortly.

2.4.4 head and tail

Sometimes we want to get access to the first few lines of a file or the last few lines of a file. The `head` and `tail` commands can be used to do that:

```

1 % head -1 words.txt
2 10th
3 % head -2 words.txt
4 10th
5 1st
6 % tail -4 words.txt
7 Zoroastrian
8 zounds
9 zucchini
10 Zurich
11 %

```

The number after the `-` specifies how many lines to print. If this option is not specified, the default is 10 lines.

`tail` also allows you to *ignore* a given number of the first lines:

```

1 % tail +4 words.txt
2 3rd
3 4th
4 5th
5 ...
6 %

```

Here, the `+4` means start printing from the 4th line, i.e. ignore the first three lines of the file. `head` and `tail` are often useful for checking the format of a file because they print out enough to see what is going on, but not so much that you lose everything else you had on the screen.

2.4.5 wc

`wc` is used to count characters, words and lines in a file:

```

1 % wc words.txt
2 24001 24001 196537 words.txt
3 % wc -l words.txt
4 24001 words.txt
5 % wc -w words.txt
6 196537 words.txt
7 %

```

Here we print the listing of lines, words and bytes, then the lines only and then the bytes (or characters) only for `words.txt`. Because this file only contains one word per line the number of each is the same for this file.

2.5 Investigating files

You should do these exercises in your `unixstuff` directory:

```
1 % cd ~/unixstuff
```

2.5.1 grep

`grep` prints lines of the input matching a given expression:

```
1 % grep 'uu' words.txt
2 continuum
3 residuum
4 vacuum
5 % grep '^x' words.txt
6 x
7 x's
8 xenon
9 xenophobia
10 xerography
11 ...
12 % grep -i '^x' words.txt
13 x
14 x's
15 Xavier
16 xenon
17 xenophobia
18 ...
19 %
```

`grep` is case sensitive by default, so we have to use the `-i` flag to set case insensitivity. Other useful flags are `-v` (line that do not match) and `-c` (produce a count only). If multiple files are given then the results are given by file.

More complex expressions can be used for the pattern, including wildcard characters. These include `*` (0 or more of the preceding character), `+` (1 or more of the preceding character), `?` (0 or 1 of the preceding character), `.` (any character), `[abc]` (any 1 of the included characters), `[^abc]` (any 1 character, except those included). When used (or if a space is used), the expression must be enclosed in single quotes.

These complex patterns are called *regular expressions* and are a powerful way of describing strings. For example

```
1 % grep '^Z.us' words.txt
2 Zeus
3 %
```

This searches for any line with a string starting with 'Z' followed by any character, followed by a 'u' character, and finished with an 's'. Regular expressions are covered in more detail in the Advanced Unix notes.

2.5.2 Redirection

Most processes initiated by Unix commands write to the *standard output* (or `stdout`) i.e. they write to the terminal window, and many take their input from the *standard input* (or `stdin`), i.e. they read it from the keyboard. There is also the *standard error* (or `stderr`), where processes write their error messages, which by default, is also to the terminal window.

Here we see the other use of `cat`: taking input from the standard input (keyboard) and printing it to the standard output (the screen). You need to type in the bits in `??normal text` and then press `<Ctrl-D>` at the end:

```
1 % cat
2 foo
3 foo
4 bar
5 bar
```



```
6 <Ctrl-D>
7 %
```

Redirection can be used to change the standard input, output and error of a program to a file instead. Running `cat` again, and redirecting standard output to a file looks like this:

```
1 % cat > list1
2 apple
3 pear
4 <Ctrl-D>
5 % cat list1
6 apple
7 pear
8 %
```

Here we have redirected the standard output to the file `list1` using the `>` operator. This will create or overwrite the file specified after the redirection operator. We can use the `>>` operator to *append* the output. If the file does not exist it acts the same as the `>` operator.

```
1 % cat >> list1
2 banana
3 orange
4 <Ctrl-D>
5 % cat list1
6 apple
7 pear
8 banana
9 orange
10 % cat >> list2
11 foo
12 bar
13 <Ctrl-D>
14 % cat list2
15 foo
16 bar
17 % cat > list2
18 beef
19 lamb
20 <Ctrl-D>
21 % cat list2
22 beef
23 lamb
24 %
```

You can see in this example above that failing to use `>>` for the second lot of output for `list2` causes the original contents (`foo` and `bar`) to be deleted.

We can also use `cat` to join (or *concatenate* files):

```
1 % cat list1 list2 > biglist
2 % cat biglist
3 apple
4 pear
5 banana
6 orange
7 beef
8 lamb
9 %
```

2.5.3 sort

`sort` can be used to sort the contents of a file:

```
1 % sort biglist
2 apple
```

```

3 banana
4 beef
5 lamb
6 orange
7 pear
8 % sort < biglist > sortedlist
9 % cat sortedlist
10 apple
11 banana
12 beef
13 lamb
14 orange
15 pear
16 %

```

Sort by default takes input from the standard input and prints the sorted contents to standard output. However, if you specify a filename, here **biglist**, it will sort that file rather than standard input. Alternatively, we can use the < operator to redirect standard input to be the file **biglist**. Here we use both operators to take input from a file and send the output to a file. Note that sort can take a series of files as input and sort these.

2.5.4 pipes

Suppose we want to run some input through several programs, such as counting the results of **grep** using **wc**. If we use redirection, it requires the creation of temporary files.

```

1 % grep 'b' biglist > filteredlist
2 % cat filteredlist
3 banana
4 beef
5 lamb
6 % wc filteredlist
7      3      3      17 filteredlist
8 %

```

This requires us to then cleanup the file when complete. Instead we can use the | to create a pipe between the two programs.

```

1 % grep 'b' biglist | wc
2      3      3      17
3 %

```

This pipe makes the output of the first program the input of the second program. These can then be chained together in arbitrarily long sequences.

```

1 % grep 'b' biglist | wc | wc
2      1      3      25
3 %

```

2.5.5 uniq

uniq merges together duplicate lines *that are adjacent to each other*. It is often used in conjunction with sort.

```

1 % cat > duplicates
2 this
3 test
4 is
5 is
6 is
7 a
8 test
9 test
10 % uniq duplicates
11 this
12 test

```

```

13 is
14 a
15 test
16 %

```

Notice that there are still two instances of `test` in the output. That's because the first `test` is separated from the other instances and so they won't be merged. If you want to remove *all* of the duplicates, sort the file first:

```

1 % sort duplicates | uniq
2 a
3 is
4 test
5 this
6 %

```

`uniq` can also count the number of times each line has occurred:

```

1 % sort duplicates | uniq -c
2 1 a
3 3 is
4 3 test
5 1 this
6 %

```

`uniq` also allows you to only print out the *unique* lines or only the *duplicate lines*:

```

1 % sort duplicates | uniq -u
2 a
3 this
4 % sort duplicates | uniq -d
5 is
6 test
7 %

```

2.6 File Access Rights

2.6.1 `ls`

Using the `-l` flag for `ls`, we can generate a *long listing* of a file or directory:

```

1 % ls -l
2 total 40
3 lrwxrwxrwx 1 tara astrop 29 Mar 6 11:10 README -> /local/usr/lib/newuser/README
4 -rw-r--r-- 1 tara astrop 35 Mar 6 10:57 biglist
5 -rw-r--r-- 1 tara astrop 17 Mar 6 10:58 filteredlist
6 -rw-r--r-- 1 tara astrop 26 Mar 6 11:09 list1
7 -rw-r--r-- 1 tara astrop 10 Mar 6 11:10 list2
8 drwx----- 3 tara astrop 4096 Mar 6 01:50 unixstuff
9 %

```

Reading left to right, the properties listed are type and permissions, number of links, owner (username), owner (group), file size, date last modified and name. The type is the first character of the listing. These have the value `-` for files, `d` for directories and `l` for symbolic links. The next nine characters are the permissions. They are in three character blocks for *user*, *group* and *world* (or other) permissions. The each block is made up of three flags: `r`, `w`, and `x`, which have slightly different meanings for files and directories. If the permission is not granted a `-` appears.

2.6.2 Access Rights for Files

`r` indicates read permission (or otherwise), that is, the presence or absence of permission to read and copy the file;

`w` indicates write permission (or otherwise), that is, the permission (or otherwise) to change a file, including to delete it;

x indicates execution permission (or otherwise), that is, the permission to execute a file, where appropriate.

2.6.3 Access Rights for Directories

r allows users to list files in the directory;

w means that users may delete files from the directory or move files into it;

x means the right to access files in the directory. This implies that you may read files in the directory provided you have read permission on the individual files.

For the above, everyone can read all the files listed, but only tara can write (or delete) them. Everyone can see into the directory `unixstuff` and access files within it (to which they have permission), but, again, only tara can modify it.

2.6.4 `chmod`

The owner of a file can change the permissions on a file using `chmod`. This allows him to hide or reveal files to other users on the computer. `chmod` takes options as follows:

u user	r read	+ add permission
g group	w write (and delete)	- remove permission
o other (world)	x execute (and access directory)	
a all		

which are used as follows:

```

1 % ls -l biglist
2 -rw-r----- 1 tara  astrop      35 Mar  6 10:57 biglist
3 % chmod g+rwk biglist
4 % ls -l biglist
5 -rw-rwx---  1 tara  astrop      35 Mar  6 10:57 biglist
6 % chmod go-rwx biglist
7 % ls -l biglist
8 -rw-----  1 tara  astrop      35 Mar  6 10:57 biglist
9 % chmod a+r,a-w biglist
10 % ls -l biglist
11 -r--r--r--  1 tara  astrop      35 Mar  6 10:57 biglist
12 % cat > biglist
13 -bash: biglist: Permission denied
14 %

```

2.7 Processes and Job Control

2.7.1 `ps`

Each process executed on the system is given a unique identifier, its PID. `ps` can be used to list your (and others') running processes. To see this we can use the `sleep` command, which will wait (or sleeps) for the given number of seconds. To be able to see this working, we also use the background modifier `&` (the ampersand). Any process executed with this at the end of its command will be run in the *background*. This allows the user to continue to type commands into the shell while the program is running.

When a process is run in the background, the shell will tell you the *jobspec* (in brackets) and the PID of the process so it can later be brought to the foreground using the `fg` command (using the last backgrounded process or taking the jobspec as its argument).

```

1 % sleep 10 &
2 [1] 20746
3 % ps
4  PID TTY          TIME CMD
5 18116 pts/4    00:00:00 bash

```

```

6 20746 pts/4    00:00:00 sleep
7 20747 pts/4    00:00:00 ps
8 %

```

Notice, there are three of your processes running (from this shell window). One of them is the shell itself (**bash**), another is the **ps** itself and the third is the **sleep** command.

2.7.2 bg and fg

A currently running process can be *stopped* without destroying it by using pressing **<Ctrl-Z>**. This sends the process to sleep and prints the jobspec to screen. **bg** can then be used to start the process again in the background. It will then run as if it had been run from the command line using the **&** command. Using **<Ctrl-Z>** is a good trick when you forgot to put an ampersand on the end of your command line. **fg** can be used to run the process in the foreground.

Here is an example of these tricks in action:

```

1 % sleep 30
2 <Ctrl-Z>
3 [1]+  Stopped                  sleep 30
4 % bg
5 [1]+  sleep 30 &
6 % fg 1
7 sleep 30
8
9 %

```

If you actually want to kill a program on the other hand, press **<Ctrl-C>**. Many people confuse the two and run **<Ctrl-Z>** and so rather than killing their programs, they just put them to sleep, so they end up with lots of sleeping processes hanging around the system.

2.7.3 jobs

This is used to list the current sleeping and backgrounded jobs:

```

1 % sleep 30 &
2 [1] 21119
3 % sleep 30 &
4 [2] 21120
5 % sleep 30
6 <Ctrl-Z>
7 [3]+  Stopped                  sleep 30
8 % jobs
9 [1]  Running                  sleep 30 &
10 [2]-  Running                  sleep 30 &
11 [3]+  Stopped                  sleep 30
12 %

```

2.7.4 kill

This is used to send *signals* to programs to stop, halt or terminate them. It takes a PID or jobspec (preceded by a **%**). **<Ctrl-C>** is the equivalent to killing on a running process.

```

1 % sleep 30
2 <Ctrl-C>
3 % sleep 30 &
4 [1] 21158
5 % jobs
6 [1]+  Running                  sleep 30 &
7 % kill %1
8 % jobs
9 [1]+  Terminated              sleep 30
10 % sleep 30
11 [1] 21165
12 % kill 21165

```

```

13 % jobs
14 [1]+  Terminated          sleep 30
15 %

```

Sometimes a specific signal sent by kill must be specified. The most common of these is `-9`, which will cause a program to exit immediately. The normal kill signal gives the program the opportunity to cleanup after itself. This does not. `-9` is most often used when a program refuses to stop.

It is also possible to kill all of your processes, in all of your shell windows at once using `kill -9 -1`. This kills everything at once, including your shell processes, so it logs you out as well.

2.8 Working with Columns of Data

There are several Unix commands that allow you to process text files that consist of columns of data. Typically these columns are separated by *whitespace*, that is, either a space or tab character⁴.

2.8.1 paste

The `paste` command takes two or more files and pastes the contents of these files together column-wise, with each column separated by a *delimiter* character. This delimiter is tab by default:

```

1 % cat > names
2 kilo
3 mega
4 giga
5 tera
6 peta
7 exa
8 % cat > powers10
9 3
10 6
11 9
12 12
13 15
14 18
15 % paste names powers10
16 kilo 3
17 mega 6
18 giga 9
19 tera 12
20 peta 15
21 exa 18
22 %

```

Note, the `kilo` and `3` etc are separated by tabs. We can change that by giving `paste` the `-d<delim>` delimiter option (or `-d <delim>`), where `<delim>` is one or more characters to use as the delimiter:

```

1 % paste -d ' >' names names powers10
2 kilo kilo>3
3 mega mega>6
4 giga giga>9
5 tera tera>12
6 peta peta>15
7 exa exa>18
8 %

```

This example demonstrates:

- you can repeat a filename
- specifying multiple delimiters
- using single quotes to stop the shell interpreting a string

⁴newlines are also counted as whitespace but not in this context because newline characters separate rows.

Multiple delimiters used in the order they are given to separate the columns (if there are more columns than delimiters they are used in cyclical order). The generalisation of this is that if you only specify one delimiter character, then it will be used to separate all columns!

Quoting arguments to avoid shell interpretation is required frequently. In our previous example, we wanted to use space as the first separator and > as the second. If we did not quote the argument, then the shell would ignore the whitespace and interpret the > as indicating redirection of standard output. If you are not sure whether the shell might be accidentally interpreting an argument, single quotes can be used to protect it from the shell.

paste can also take redirected standard input where the placeholder - means read the standard input:

```
1 % paste -d '_' - powers10 < names
2 kilo_3
3 mega_6
4 giga_9
5 tera_12
6 peta_15
7 exa_18
8 %
```

Here, **names** has been *redirected* as standard input in place of the keyboard input. **paste** then reads from standard input when the - placeholder is used.

The final use of **paste** is to turn a flat file into multiple columns (something we only discovered reading the manpages recently, but is a very useful feature!). This involves repeated use of the - placeholder filename and is a very cunning trick:

```
1 % paste - - - < names
2 kilo mega giga
3 tera peta exa
4 %
```

2.8.2 cut

Not surprisingly, there is a command to do the reverse of pasting, and that command is called **cut**. This allows you to select specific columns from a given file or standard input.

Firstly, have a look at the **columns.txt** file using **less** or **cat**. Notice that the entries are separated in columns:

```
1 % cat columns.txt
2 15:48:55.81 -40:32:17.5 J1600M40 2 1393.1 333.9 3
3 15:49:10.47 -40:14:18.7 J1600M40 2 1380.9 397.2 3
4 15:49:12.16 -40:02:13.1 J1600M40 2 1381.1 439.6 3
5 ...
6 %
```

The columns are Right Ascension, Declination, Mosaic Name and various other source information.

So, to get a list of just the declinations we can go:

```
1 % cut -d ' ' -f 2 columns.txt
2 -40:32:17.5
3 -40:14:18.7
4 -40:02:13.1
5 ...
6 %
```

The **-d** option allows you to specify the column delimiter. Try using **-d ':'** instead.

Cut allows you to specify the fields in a fairly flexible manner with the **-f** option:

-f 1,3,5	selects columns 1, 3 and 5 (surprising huh!)
-f 4-6	selects columns 4, 5 and 6
-f -3	selects columns 1, 2 and 3
-f 4-	selects columns 4 and any more columns to the right

2.8.3 `awk` and `gawk`

AWK is a programming language for processing text files containing columns of data. It was developed by Aho, Weinberger and Kernighan more than 20 years ago. AWK partly inspired the development of *Perl*, which is a very popular modern scripting language.

`awk` (or the Gnu version `gawk`) is the Unix command for interpreting AWK programs on text input. AWK allows you to write quite complicated expressions for calculations involving columns of data. We are only going to cover a few simple but powerful ones here.

The first is selecting rows based on properties of particular columns:

```
1 % paste names powers10 > prefixes
2 % awk '$2 > 6' prefixes
3 giga 9
4 tera 12
5 peta 15
6 exa 18
7 % awk '$1 == "giga"' prefixes
8 giga 9
```

The first example only prints the row out if column 2 (represented by `$2`) has a numerical value larger than 6. The second example only prints out rows that have the word `giga` in their first column. There are lots of powerful matching techniques which we will discuss in later labs and lectures.

```
1 % awk '{ print $2, $1; }' prefixes
2 3 kilo
3 6 mega
4 9 giga
5 12 tera
6 15 peta
7 18 exa
8 % awk '{ print $1, NR*10; }' prefixes
9 kilo 10
10 mega 20
11 giga 30
12 tera 40
13 peta 50
14 exa 60
15 %
```

It is also possible to produce different output, rather than simply selecting interesting rows. In the first example we swap the two columns around by printing `$2` before `$1`. Notice that the `print` statement needs to be within *braces* (or squiggly brackets) and followed by a semi-colon.

The second example demonstrates that we can also calculate new values on the fly and print these out as well. Here we have calculate the powers of 2 for SI units rather than the powers of 10. We use a special variable `NR` which holds the line number we are up to in the input (the *number of records*). Another useful special variable is `NF` which is the number of fields in the current row.

In both of these examples, the output columns were separated by spaces rather than tabs, because this is the default output field separator (`OFS`) which we can change:

```
1 % awk 'BEGIN { OFS=">"; } { print $1, $2; }' prefixes
2 kilo>3
3 mega>6
4 giga>9
5 tera>12
6 peta>15
7 exa>18
```

Here we use the special `BEGIN` rule to set the value of the output field separator *before* AWK starts processing any lines of input. We can also use `BEGIN` to set the input field separator (`FS`) before processing begins.

There is also an equivalent `END` rule which is run after the input has been processed:


```
1 % awk '{ sum = sum + $2; } END { print sum; }' prefixes
2 63
```

This program stores the sum of all of the values in column 2 in a variable called `sum`, and then after it has processed all of the rows, prints out the value of `sum`.

AWK has all of the power of a programming language rather than a simple command line tool. You can find out more about what it can do from the manpage `awk (1)`, but it is a rather terse description.

2.8.4 More on `sort`

`sort` can also work in a column aware manner. By default it sorts on the first column of input in ascending alphabetical order. However, all of these can be changed with command line arguments.

Firstly, you can reverse the sorting order with the `-r` option:

```
1 % sort -r prefixes
2  tera  12
3  peta  15
4  mega   6
5  kilo   3
6  giga   9
7  exa   18
8  %
```

Secondly, you can sort on a different column using the `-k` option to indicate which column to sort on first:

```
1 % sort -k 2 prefixes
2  tera  12
3  peta  15
4  exa   18
5  kilo   3
6  mega   6
7  giga   9
8  %
```

Hmm, what went wrong here? The rows don't seem to be sorted in any sensible order at all? Well in fact they are sorted, but the sorted order is still *lexicographic*, that is, phone book ordering in alphanumeric ordering. In lexicographic ordering, the length of the word is not important, only the relative order of the first characters in the word. For example, `aardvark` comes before `ant` because 'a' comes before 'n'.

Therefore, when we treat numbers in a lexicographic way, we end up putting all of the numbers starting with the digit '1' first. If we want to treat the column as numeric, we can use the `-n` flag (in first column sorting) or add `n` to the column specifier with the `-k` option.

```
1 % sort -k 2n prefixes
2  kilo   3
3  mega   6
4  giga   9
5  tera  12
6  peta  15
7  exa   18
8 % sort -k 2rn prefixes
9  exa   18
10 peta  15
11 tera  12
12 giga   9
13 mega   6
14 kilo   3
15 %
```

The Gnu version also provides a `-g` option (or `g` for use with the `-k` option) which treats the sorting column as a floating point number with the exponent notation.

2.9 Other Useful Commands

2.9.1 `exit`

This ends the current interactive shell session and so causes you to be logged out. If you still have processes running, it will hang there until the processes have finished.

2.9.2 `gzip` and `gunzip`

`gzip` and `gunzip` are the Gnu zipping programs for compressing and decompressing files. Notice that they remove the original and compressed version of the files after they have run:

```
1 % gzip biglist
2 % ls biglist*
3 biglist.gz
4 % gunzip biglist.gz
5 % ls biglist*
6 biglist
7 %
```

2.9.3 `tar`

`gzip` is designed to compress a single file, so how do you compress and transfer multiple files around the system?

The `tar` *archiving* program collates many files and directories into a single file. The `c` flag specifies create an archive file, the `x` flag specifies extract an archive and the `t` flag specifies list the contents of an archive without extracting files. The `v` flag specifies verbose output and the `f` flag indicates that the following argument is the output file.

The `z` flag can be added to use `gzip` compression and decompression, which is the same as running `gzip/gunzip` after/before running the `tar` command.

```
1 % tar cvf lists.tar list1 list2
2 list1
3 list2
4 % tar xvf lists.tar
5 list1
6 list2
7 % tar cvzf lists.tar.gz list1 list2
8 list1
9 list2
10 % ls -l lists.tar*
11 -rw-rw-rw-  1 info1903  cs1      10240 Mar  6 14:56 lists.tar
12 -rw-rw-rw-  1 info1903  cs1         178 Mar  6 14:57 lists.tar.gz
13 % tar tzvf
```

2.9.4 Tab completion

The tab can be used to auto-complete shell strings on the command line. In its simplest form it allows for completion of filenames. Later versions of Bash also allow completion of arguments and restrict irrelevant file types. Try typing a single character and then pressing `<tab>`. If there is only one possible completion, the shell will automatically complete it for you. If there is more than one completion, you will hear a beep instead. Press `<tab>` again to see the list of possible completions.

2.9.5 Wildcards

Wildcards can be used on the command line to identify files. `*` will match zero or more of any character. `?` will find exactly one of any character. Square brackets can be used to create sets of characters to match from. `[abc]` will match either the 'a', 'b' or 'c' in a filename.

```
1 % ls list*
2 list1
3 list2
4 % ls *list
5 biglist
```

```

6 filteredlist
7 % ls *list*
8 biglist
9 filteredlist
10 list1
11 list2
12 % ls ?list
13 biglist
14 %

```

2.9.6 Shell Variables

Variables are a way of passing information from the shell to programs when you run them. Programs look “in the environment” for particular variables and if they are found will use the values stored in them. Some are set by the system, others by you, yet others by the shell, or any program that loads another program.

Standard Unix variables are split into two categories, environment variables and shell variables. In broad terms, shell variables apply only to the current instance of the shell and are used to set short-term working conditions; environment variables have a farther reaching significance, and those set at login are valid for the duration of the session. By convention, environment variables have UPPER CASE and shell variables have lower case names.

2.9.7 echo

The **echo** command writes its arguments to standard output and can therefore be used to display the contents of variables. It will repeat to standard output whatever followed it on the command line. Environment and shell variables can be accessed by using the **\$** symbol before their name.

```

1 % echo $USER
2 tara
3 %

```

Some common environment variables are:

- OSTYPE the operating system type
- USER your login name
- HOME the path name of your home directory
- HOST the name of the computer you are using
- ARCH the architecture of the computers processor
- DISPLAY the name of the computer screen to display X windows
- PRINTER the default printer to send print jobs
- PATH the directories the shell should search to find a command

2.9.8 printenv

printenv can be used to display all the currently set environment variables.

2.10 Acknowledgements

These notes are based on those used for our course INFO1903 <http://www.it.usyd.edu.au/~info1903>. Thanks to James Gorman for compiling the first version of the notes. They are largely derived from the notes written by M Stonebank from the University of Surrey. <http://www.ee.surrey.ac.uk/Teaching/Unix/> Material was also taken from the University of Utah's Unix command summary at <http://www.math.utah.edu/computing/unix/unix-commands.html>