

# Unix Systems: Shell Scripting (II)

Bruce Beckles  
University of Cambridge Computing Service



# Introduction

- Who:
  - Bruce Beckles, e-Science Specialist, UCS
- What:
  - Unix Systems: Shell Scripting (II) course
  - Follows on from “[Unix Systems: Shell Scripting \(I\)](#)”
  - Part of the *Scientific Computing* series of courses
- Contact (questions, etc):
  - [escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)
- Health & Safety, etc:
  - Fire exits
- **Please switch off mobile phones!**

[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

2

As this course is part of the Scientific Computing series of courses run by the Computing Service, all the examples that we use will be more relevant to scientific computing than to system administration, etc.

This does not mean that people who wish to learn shell scripting for system administration and other such tasks will get nothing from this course, as the techniques and underlying knowledge taught are applicable to shell scripts written for almost any purpose. However, such individuals should be aware that this course was not designed with them in mind.

For details of the “Unix Systems: Shell Scripting (I)” course, see:  
<http://www.cam.ac.uk/cs/courses/coursedescript1>

## What we don't cover

- Different types of shell:
  - We are using the [Bourne-Again Shell](#) (bash).
- Differences between versions of bash
- Very advanced shell scripting – try these courses instead:
  - “[Unix Systems: Shell Scripting \(III\)](#)”
  - “[Programming: Python for Absolute Beginners](#)”

`bash` is probably the most common shell on modern Unix/Linux systems – in fact, on most modern Linux distributions it will be the default shell (the shell users get if they don't specify a different one). Its home page on the WWW is at:

<http://www.gnu.org/software/bash/>

We will be using `bash` 3.0 in this course, but everything we do should work in `bash` 2.05 and later. Version 3.0 and version 2.05 (or 2.05a or 2.05b) are the versions of `bash` in most widespread use at present. Most recent Linux distributions will have one of these versions of `bash` as one of their standard packages. The latest version of `bash` (at the time of writing) is `bash` 3.2, which was released on 12 October, 2006.

For details of the “Unix Systems: Shell Scripting (III)” course, see:

<http://www.cam.ac.uk/cs/courses/coursedescriptwkshp>

For details of the “Programming: Python for Absolute Beginners” course, see:

<http://www.cam.ac.uk/cs/courses/coursedescriptprog.html#python>

# Outline of Course

1. Prerequisites & recap of “Shell Scripting (I)” course
2. Shell functions

*SHORT BREAK*

3. Command substitution
4. The `mktmp` command
5. Handling data from standard input
  - Reading values from standard input
  - Pipelines
  - Loop constructs: `while`

*SHORT BREAK*

6. More `while` loops:
  - Shell arithmetic
  - Tests

Exercise (~16:30)

[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

4

The course officially finishes at 17.00, but the intention is that the lectured part of the course will be finished by about 16.30 and the remaining time is for you to attempt an exercise that will be provided. If you need to leave before 17.00 (or even before 16.30), please do so, but don't expect the course to have finished before then. If you do have to leave early, please leave quietly and ***please make sure that you fill in a green Course Review form*** and leave it at the front of the class for collection by the course giver.

# Follow-on course

## Unix Systems: Shell Scripting (III):

- Better, more robust (**handle errors!**) shell scripts

**We *strongly* encourage you to attend this follow-on course.**

For details of the “Unix Systems: Shell Scripting (III)” course, see:

<http://www.cam.ac.uk/cs/courses/coursedescribe/linux.html#scriptwkshp>

# Pre-requisites

- Ability to use a text editor under Unix/Linux:
  - Try gedit if you aren't familiar with any other Unix/Linux text editors
- Familiarity with the Unix/Linux command line (“[Unix System: Introduction](#)” course)
- Familiarity with material covered in “[Unix Systems: Shell Scripting \(I\)](#)” course:
  - Simple shell scripts: linear lists of commands
  - Simple use of shell variables and parameters
  - Simple command line processing
  - Output redirection
  - for loops

[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

6

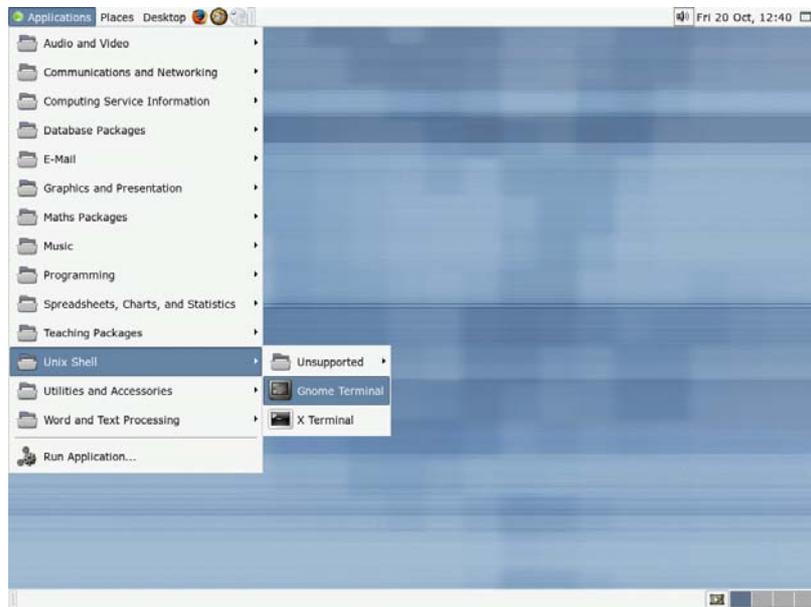
For details of the “Unix System: Introduction” course, see:

<http://www.cam.ac.uk/cs/courses/coursedescript/linux.html#unix>

For details of the “Unix Systems: Shell Scripting (I)” course, see:

<http://www.cam.ac.uk/cs/courses/coursedescript/linux.html#script1>

# Start a shell

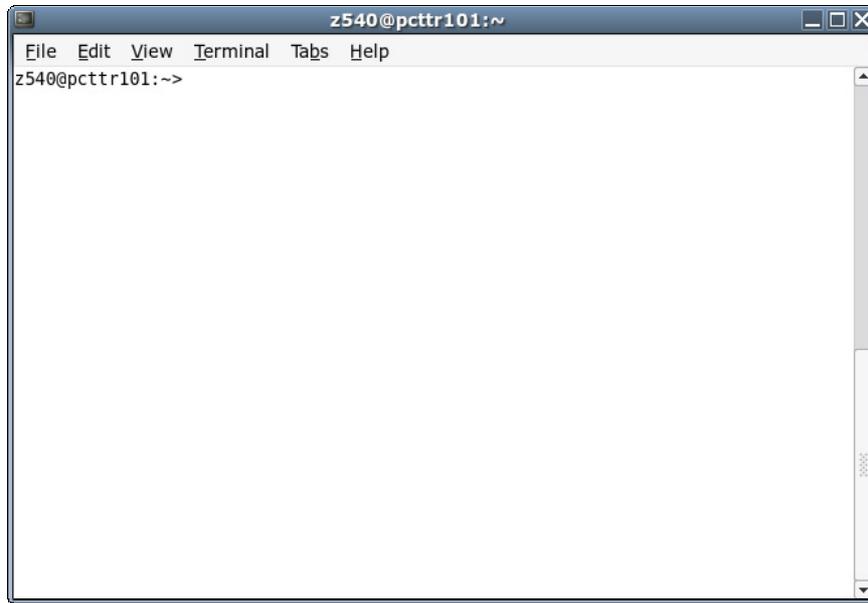


[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

7

# Screenshot of newly started shell



## Recap: What is a shell script?

- **Text** file containing commands understood by the shell
- Very **first** line is special:  
`#!/bin/bash`
- File has its **executable** bit set  
`chmod +x`

Recall that the `chmod` command changes the permissions on a file. `chmod +x` sets the executable bit on a file, i.e. it grants permission to execute the file. Unix file permissions were covered in the “Unix System: Introduction” course, see:

<http://www.cam.ac.uk/cs/courses/coursedescript/linux.html#unix>

## Sample program: `iterator`

```
> ./iterator 100 100 1000 0.05
```

```
x dimension of grid:    100  
y dimension of grid:    100  
Number of iterations:  1000  
Epsilon:                0.050000
```

```
Output file:           output.dat
```

```
Iterations took 2.100 seconds
```

The `iterator` program is in your home directory. It is a program written specially for this course, but we'll be using it as an example program for pretty general tasks you might want to do with many different programs. Think of `iterator` as just some program that takes some input on the command line and then produces some output (on the screen, or in one or more files, or both), e.g. a scientific simulation or data analysis program.

The `iterator` program takes 4 numeric arguments on the command line: 3 positive integers and 1 floating-point number. It always writes its output to a file called `output.dat` in the current working directory, and also writes some informational messages to the screen.

The `iterator` program is not as well behaved as we might like (which, sadly, is also typical of many programs you will run). The particular way that `iterator` is not well behaved is this: every time it runs it creates a file called `running` in the current directory, and it will not run if this file is already there (because it thinks that means it is already running). Unfortunately, it doesn't remove this file when it has finished running, so we have to do it manually if we want to run it multiple times in the same directory.

# Writing simple shell scripts

Suppose we want to run `iterator` several times, keeping the output file from each run. We might do it like this:

1. Delete any file called `running`  
> **rm running**
2. Run `iterator` with some parameters  
> **./iterator 100 100 1000 0.05**
3. Rename `output.dat`  
> **mv output.dat output-0.05.dat**
4. Repeat the above steps for the following parameter sets:  
    100 100 1000 0.1  
    100 100 1000 0.15

We could easily write a very simple shell script that does the above task. Basically, we want to run the `iterator` program three times with a different parameter set each time. Note that only the last parameter changes between each run, and that is the parameter we insert into the output file name when we rename it to stop it being overwritten by the next run.

Those of you who did the “Shell Scripting (I)” course will have constructed just such a shell script as an exercise in that course.

# A simple shell script

```
#!/bin/bash

# Remove left over running file
rm -f running
# Run iterator
./iterator 100 100 1000 0.05
# Rename output file
mv output.dat output-0.05.dat

# Run iterator again
rm -f running
./iterator 100 100 1000 0.1
mv output.dat output-0.1.dat

# Run iterator yet again
rm -f running
./iterator 100 100 1000 0.15
mv output.dat output-0.15.dat
```

A version of this shell script called `run-iterator.sh`, with more extensive comments than those shown on this slide, is in your home directory.

## Recap: Very simple shell scripts

- Linear lists of commands
- Just the commands you'd type interactively put into a file
- Simplest shell scripts you'll write

# Shell variables and parameters

Shell variables hold data, much like variables in a program:

```
> VAR="My variable"
> echo "${VAR}"
My variable
```

Shell parameters are special variables set by the shell:

- Positional parameter 0 holds the name of the shell script
- Positional parameter 1 holds the first argument passed to the script; positional parameter 2 holds the second argument passed to the script, etc
- Special parameter @ expands to values of all positional parameters (starting from 1)
- Special parameter # expands to the number of positional parameters (not including 0)

We create shell variables by simply assigning them a value (as above for the shell variable `VAR`). We can access the value of a shell variable using the construct `${VARIABLE}` where `VARIABLE` is the name of the shell variable. Note that there are **no** spaces between the name of the variable, the equal sign (=) and the variable's value in double quotes. *This is very important* as *whitespace* (spaces, tabs, etc) is significant in the names and values of shell variables.

Also note that although we can assign the value of one shell variable to another shell variable, e.g. `VAR1="${VAR}"`, the two shell variables are in fact completely separate from each other, i.e. each shell variable can be changed independently of the other. Changing the value of one will not affect the other. So `VAR1` (in this example) is *not* a "pointer" to or an "alias" for `VAR`.

*Shell parameters* are special variables set by the shell. Many of them cannot be modified, or cannot be directly modified, by the user or by a shell script. Amongst the most important parameters are the *positional parameters* and the other shell parameters associated with them.

The positional parameters are set to the arguments that were given to the shell script when it was started, with the exception of positional parameter 0, which is set to the name of the shell script. So, if `myscript.sh` is a shell script, and I ran it by typing:

```
./myscript.sh argon hydrogen mercury
```

then positional parameter

- 0 = `./myscript.sh`
- 1 = `argon`
- 2 = `hydrogen`
- 3 = `mercury`

and all the other positional parameters are not set.

The special parameter `@` is set to the value of all the positional parameters, starting from the first parameter, passed to the shell script, each value being separated from the previous one by a space. You access the value of this parameter using the construct `${@}`. If you access it in double quotes – as in `"${@}"` – then the shell will treat each of the positional parameters as a separate *word* (which is what you normally want).

The special parameter `#` is set to the number of positional parameters *not counting positional parameter 0*. Thus it is set to the number of arguments passed to the shell script, i.e. the number of arguments on the command line when the shell script was run.

# Shell parameters

- Positional parameters (`${0}`, `${1}`, etc)
- Value of all arguments passed: `${@}`
- Number of arguments: `${#}`

```
> ~/examples/params.sh 0.5 62 38 hydrogen
```

```
This script is /home/x241/examples/params.sh
```

```
There are 4 command line arguments.
```

```
The first command line argument is: 0.5
```

```
The second command line argument is: 62
```

```
The third command line argument is: 38
```

```
Command line passed to this script: 0.5 62 38 hydrogen
```

[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

15

In the `examples` subdirectory of your home directory there is a script called `params.sh`. If you run this script with some command line arguments it will illustrate how the positional parameters and related shell parameters work. Note that even if you type exactly the command line on the slide above your output will probably be different as the script will be in a different place for each user.

The positional parameter `0` is the name of the shell script (it is the name of the command that was given to execute the shell script).

The positional parameter `1` contains the first argument passed to the shell script, the positional parameter `2` contains the second argument passed and so on.

The special parameter `#` contains the number of arguments that have been passed to the shell script. The special parameter `@` contains all the arguments that have been passed to the shell script.

## Redirection (>, >>)

Redirect output to a file, *overwriting* file if it exists:

```
command > file
```

Redirect output to a file, *appending* it to that file:

```
command >> file
```

The > operator redirects the output from a command to a file, **overwriting** that file if it exists. You place this operator at the end of the command, after all of its arguments. This is equivalent to using `1>filename` which means “redirect file descriptor 1 (*standard output*) to the file `filename`, **overwriting** it if it exists”.

(Unsurprisingly, `2>filename` means “redirect file descriptor 2 (*standard error*) to the file `filename`, **overwriting** it if it exists”. And it will probably come as no shock to learn that `descriptor>filename` means “redirect file descriptor `descriptor` to the file `filename`, **overwriting** it if it exists”, where *descriptor* is the number of a valid file descriptor.)

You may think that this “overwriting” behaviour is somewhat undesirable – you can make the shell refuse to overwrite a file that exists, and instead return an error, using the `set` shell builtin command as follows:

```
set -o noclobber
```

or, equivalently:

```
set -C
```

The >> operator redirects the output from a command to a file, **appending** it to that file. You place this operator at the end of the command, after all of its arguments. If the file does not exist, it will be created.

# for

Execute some commands once for each value in a collection of values

```
for VARIABLE in <collection of values> ; do
    <some commands>
done
```

Examples:

```
myCOLOURS="red green blue"
for zzVAR in ${myCOLOURS} ; do
    echo "${zzVAR}"
done

for zzVAR in * ; do
    ls -l "${zzVAR}"
done
```

[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

17

We can repeat a set of commands using a `for` loop. A `for` loop repeats a set of commands once for each element in a collection of values it has been given. We use a `for` loop like this:

```
for VARIABLE in <collection of values> ; do
    <some commands>
done
```

where *<collection of values>* is a set of one or more values (strings of characters). Each time the `for` loop is executed the shell variable *VARIABLE* is set to the next value in *<collection of values>*. The two most common ways of specifying this set of values is by putting them in a another shell variable and then using the `${}` construct to get its value (note that this should *not* be in quotation marks), or by using a wildcard (e.g. `*`) to specify a collection of file names (*pathname expansion*). *<some commands>* is a list of one or more commands to be executed.

Note that you can put the `do` on a separate line, in which case you can omit the semi-colon (`;`):

```
for VARIABLE in <collection of values>
do
    <some commands>
done
```

There are some examples of how to use it in the `for.sh` script in the `examples` directory of your home directory.

## run-once.sh

```
#!/bin/bash
# Remove left over running file
rm -f running
# Write to logfile
echo "" >> logfile
date >> logfile
echo "Running iterator with ${@}" >> logfile
# Run iterator with passed arguments
./iterator "${@}" > "stdout-${4}"
# Remove left over running file
rm -f running
# Rename output file
mv output.dat "output-${4}.dat"
# Write to logfile
echo "Output file: output-${4}.dat" >> logfile
echo "Standard output: stdout-${4}" >> logfile
```

The file `run-once.sh` is in the `scripts` directory. Its contents are shown above. It runs the `iterator` program once with a given set of parameters, keeping the screen output as well as the output file with names based on the fourth parameter it is given. It also keeps a record of what it is doing in the file `logfile` in the current directory. Making your shell scripts keep a record of what they are doing is an extremely good idea, especially if they are going to run for a long time or on a remote machine or when you are not around.

Notice that we have something written to the `logfile` *before* we start running the `iterator` program *and* something *after* it is finished. This means that if the shell script crashes or is stopped before it is finished there is a very good chance we'll be able to tell from the log file as it will not have the "Output file:" or "Standard output:" lines in it. There are better, more sophisticated ways of checking whether things have gone wrong, but this is a nice simple one that is well worth remembering.

# multi-run.sh

```
#!/bin/bash

# Parameters that stay the same each run
myFIXED_PARAMS="100 100 1000"

# Run iterator program once for each argument
# Note: *no* quotes around ${myFIXED_PARAMS}
#       or they'll be interpreted as one argument!
for zzARGS in "${@}" ; do
    ~/scripts/run-once.sh ${myFIXED_PARAMS} ${zzARGS}
done

> cd
> rm -f *.dat stdout-* logfile
> scripts/multi-run.sh 0.05 0.1 0.15 0.2 0.25 0.3 0.35 0.4 0.45 0.5
```

[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

19

The file `multi-run.sh` in the `scripts` directory (shown above) takes one or more command line arguments and then runs the `run-once.sh` script (which in turn runs the `iterator` program) with 4 arguments - 3 that are always the same and are hard-coded into the script, and 1 that is one of its command line arguments. It does this repeatedly until there are no more of its command line arguments. This script is much more versatile than the script we saw earlier that ran `iterator` 3 times with 3 different parameter sets. Modifying that script for each different set of values we might want to run would have rapidly become extremely tedious, whereas we don't need to modify this script at all - we just run it with different arguments.

Note that when we use the value of the shell variable `myFIXED_PARAMS` we *don't* surround it with quotes - if we did then it would be treated as a single value instead of as 3 separate values (when the shell treats spaces in this way - as a separator between values - it is called *word splitting*).

Give it a try - change to your home directory and type the following commands (the `rm` command is to remove the files produced by our previous runs of earlier scripts):

```
rm -f *.dat stdout-* logfile
scripts/multi-run.sh 0.05 0.1 0.15 0.2 0.25 0.3 0.35 0.4 0.45 0.5
more logfile
```

And finally do a `ls` of your home directory and see what files have been produced.

## Exercise from Shell Scripting (I)

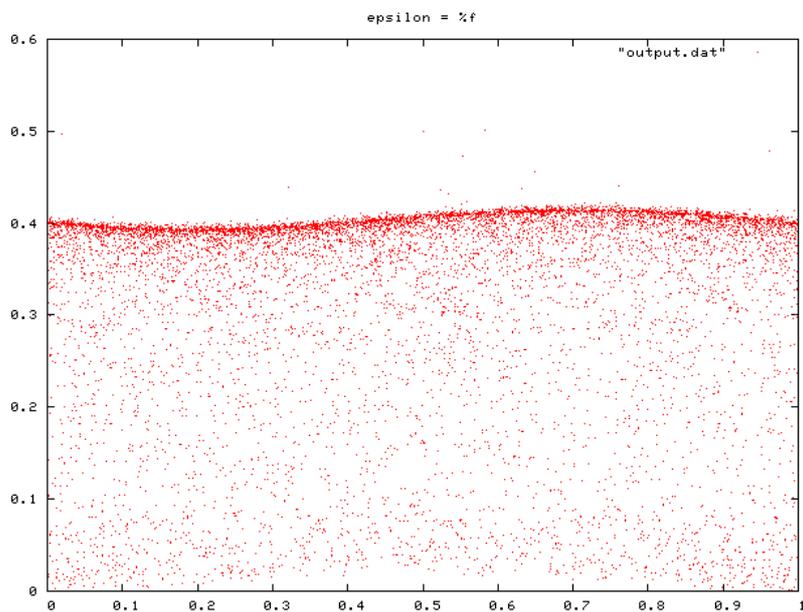
We have a directory that contains the output of several runs of the `iterator` program in separate files. We have a file of commands that will turn the output into a graph (using [gnuplot](#)). We want to write a shell script that turns the output from each run into a graph.

We are specifically using the `gnuplot` program and the output of the `iterator` program we've met before. (`gnuplot` is a program that creates graphs, histograms, etc from numeric data.) Think of this task as basically: I have some data sets and I want to process them all in the same way. My processing might produce graphical output, as here, or it might produce more data in some other format.

If you haven't met `gnuplot` before, you may wish to look at its WWW page:

<http://www.gnuplot.info/>

# Output of gnuplot



[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

21

If you want to get an idea of what we're trying to do, you can try the following:

```
> cd
> cp gnuplot/iterator.gplt .
> cp output-0.05.dat output.dat
> ls output.png
/bin/ls: output.png: No such file or directory
> gnuplot iterator.gplt
> rm output.dat
> ls output.png
output.png
> eog output.png &
```

Note that the output of "ls output.png" may look slightly different - in particular, the colours may be slightly different shades (assuming you are reading these notes in colour).

## Details of exercise

What we want to do is, *for* each output file:

1. Rename (or copy) the output file we want to process to `output.dat`  
> `mv output-0.05.dat output.dat`
2. Run `gnuplot` with the `iterator.gplt` file  
> `gnuplot iterator.gplt`
3. Rename (or delete if you copied the original output file) `output.dat`  
> `mv output.dat output-0.05.dat`
4. Rename `output.png`  
> `mv output.png output-0.05.dat.png`

The exercise set at the end of the “Unix Systems: Shell Scripting (I)” course was to create a shell script that does the above task. Basically, for each of the `.dat` files produced by the `multi-run.sh` script, the shell script should run `gnuplot` on it to create a graph (which will be stored as a `.png` file). The `iterator.gplt` file provided will only work if the `.dat` file is called `output.dat` and is in the current directory. Also, `gnuplot` should not be allowed to overwrite each `.png` file, so the shell script must rename each `.png` file after `gnuplot` has created it.

# multi-gnuplot1.sh

```
#!/bin/bash

# Run gnuplot program once for each output file
for zzFILES in output-*.dat ; do

    # Rename output file to output.dat
    mv "${zzFILES}" output.dat

    # Run gnuplot
    gnuplot iterator.gplt

    # Rename output.dat to original name
    mv output.dat "${zzFILES}"

    # Rename output.png
    mv output.png "${zzFILES}.png"

done
```

[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

23

So here's one solution to that exercise. This file (`multi-gnuplot1.sh`) is in the `gnuplot` directory.

It takes each file whose name is of the form `output-<something>.dat` (where the `<something>` can be any set of characters that can appear in a filename) in turn and renames it to `output.dat`, runs `gnuplot`, then renames the file back to its original name, and renames the `output.png` file to `output-<something>.dat.png`.

## multi-gnuplot2.sh

```
#!/bin/bash

# Run gnuplot program once for each output file
for zzFILES in output-*.dat
do
    # Copy output file to output.dat
    cp -f "${zzFILES}" output.dat

    # Run gnuplot
    gnuplot iterator.gplt

    # Delete output.dat file
    rm -f output.dat

    # Rename output.png
    mv output.png "${zzFILES}.png"
done
```

[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

24

...and here's another solution. This file (`multi-gnuplot2.sh`) is in the `gnuplot` directory.

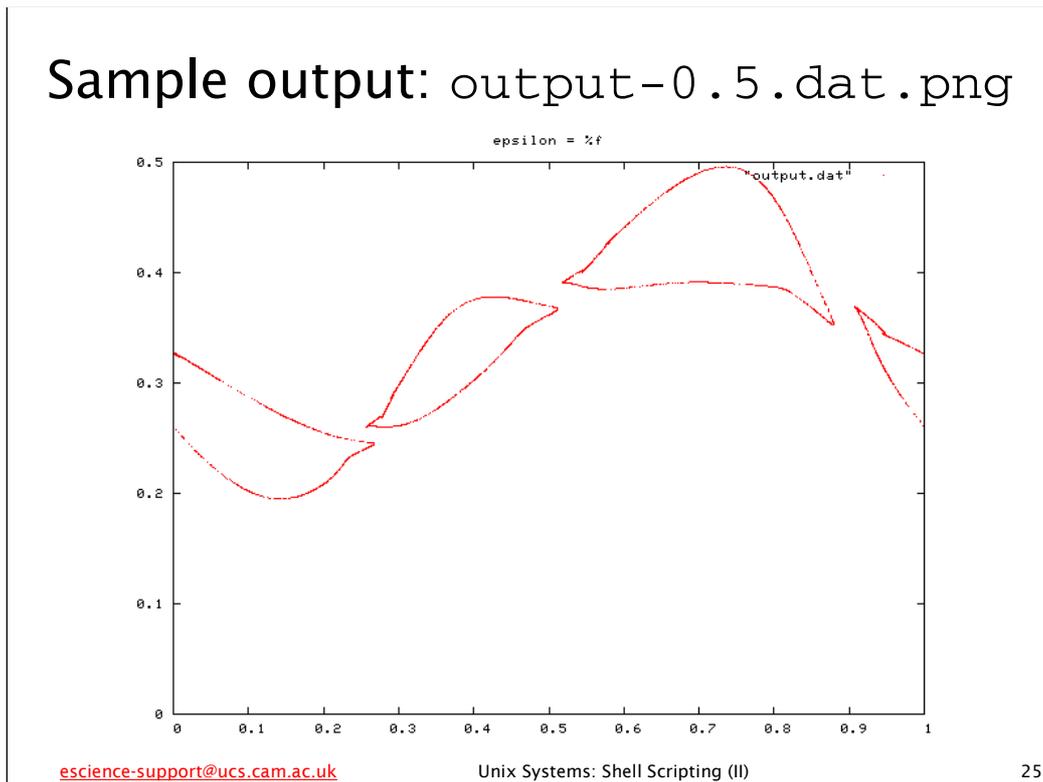
It takes each file whose name is of the form `output-<something>.dat` (where the `<something>` can be any set of characters that can appear in a filename) in turn and copies it to `output.dat`, runs `gnuplot`, then deletes the copy, and renames the `output.png` file to `output-<something>.dat.png`.

These two shell scripts are functionally equivalent - you can use whichever you like and the results will be identical.

Note that one purely cosmetic difference between them is that one has the `do` keyword on the same line as the `for` keyword (with a semi-colon `;` before the `do`) whilst the other has the `do` keyword on a separate line (and no semi-colon). Some people feel that it makes scripts more readable to put the `do` on a separate line.

However, whether you put the `do` on the same line as the `for` (and use the semi-colon) or put it on a different line is entirely a matter of style and personal preference - well, you want some outlet for your individuality, don't you? ☺

## Sample output: output-0.5.dat.png



You can try out one of these scripts if you want. First, create some output files for the script to process:

```
> cd  
> rm -f *.dat stdout-* logfile  
> scripts/multi-run.sh 0.05 0.1 0.15 0.2 0.25 0.3 0.35 0.4 0.45 0.5
```

Now, make sure that the `iterator.gplt` file is in your current directory:

```
> cp gnuplot/iterator.gplt .
```

Now run one of the scripts, either `multi-gnuplot1.sh` or `multi-gnuplot2.sh`, it doesn't matter which:

```
> scripts/multi-gnuplot1.sh
```

Now do an `ls` to see what files have been created, and then try viewing some of them:

```
> eog output-0.5.dat.png &
```

Your solutions to this exercise (you did do it, didn't you?) should have been similar to the ones presented here. If they weren't, or if you had problems with the exercise, please let the course giver or demonstrator know.

# Shell functions

```
> cd
> cat hello-function.sh
#!/bin/bash
function hello()
{
    # This is a shell function.
    echo "Hello!"
    echo "I am function ${FUNCNAME}."
}

> ./hello-function.sh
>
```

[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

26

Shell functions are similar to functions in most high-level programming languages. Essentially they are “mini-shell scripts” (or bits of shell scripts) that are invoked (*called*) by the main shell script to perform one or more tasks. When called they can be passed arguments (parameters), as we will see later, and when they are finished they return control to the statement in the shell script immediately after they were called.

To define a function, you just write the following at the start of the function:

```
function function_name()
```

```
{
```

where *function\_name* is the name of the function. Then, after the last line of the function you put a line with just a closing curly brace (}) on it:

```
}
```

Note that *unlike* function definitions in most high level languages you don’t list what parameters (arguments) the function takes. This is not so surprising when you remember that shell functions are like “mini-shell scripts” – you don’t explicitly define what arguments a shell script takes either.

Like functions in a high-level programming language, defining a shell function doesn’t actually make the shell script do anything – the function has to be called by another part of the shell script before it will actually *do* anything.

`FUNCNAME` is a special shell variable (introduced in version 2.04 of `bash`) that the shell sets within a function to the name of that function. When not within a function, the variable is unset.

# Calling a shell function

```
> gedit hello-function.sh &
#!/bin/bash
function hello()
{
    # This is a shell function.
    echo "Hello!"
    echo "I am function ${FUNCNAME}."
}

hello

> ./hello-function.sh
Hello!
I am function hello.
>
```

[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

27

Start your favourite editor (or `gedit` if you don't have a preference) and modify the file `hello-function.sh` in your home directory as shown above. *Make sure you save the file after you've modified it or your changes won't take effect.*

Now try running the shell script again:

```
> ./hello-function.sh
Hello!
I am function hello.
>
```

This time it actually does something - the function `hello` is called and does what we would expect.

You call a shell function by just giving its name (just as you would with any of the standard Unix commands (or shell builtin commands) that we've met). Note that you don't put brackets after the name of the function when you call it. You only do that when you first define the function. That's one of the ways that the shell figures out that you are trying to define a shell function.

## Shell function arguments (1)

```
> gedit hello-function.sh &
#!/bin/bash
function hello()
{
    # This is a shell function.
    echo "Hello, ${1}"
    echo "I am function ${FUNCNAME}."
}

hello

> ./hello-function.sh Dave
Hello,
I am function hello.
>
```

[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

28

Modify the file `hello-function.sh` in your home directory as shown above. *Make sure you save the file after you've modified it or your changes won't take effect.*

Recall that the positional parameter `1` (whose value is accessed using the construct `${1}`) contains the value of the first argument passed to the shell script (or is unset if no arguments are passed). So what would we expect the above shell script to do? Surely, it will print out "Hello, *<whatever argument we gave it>*"?

(For the pedants amongst you: *<whatever argument we gave it>* means whatever argument we passed the shell script on the command line when we invoked it - "Dave" in the above example.)

Apparently not. Maybe something's wrong with our shell script? Maybe positional parameter `1` isn't being set correctly? Let's try some debugging and see.

## Shell function arguments (2)

```
> gedit hello-function.sh &
#!/bin/bash
function hello()
{
    # This is a shell function.
    echo "Hello, ${1}"
    echo "I am function ${FUNCNAME}."
}

echo "First argument: ${1}"
hello

> ./hello-function.sh Dave
First argument: Dave
Hello,
I am function hello.
```

[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk) Unix Systems: Shell Scripting (II) 29

Modify the file `hello-function.sh` in your home directory as shown above. *Make sure you save the file after you've modified it or your changes won't take effect.*

This is a simple but useful debugging trick for shell scripts. When something isn't working right, make the shell script print out the values of all the shell variables, environment variables or shell parameters that you are interested in *just before* the point where you think it is going wrong.

In this case, what this shows us is that positional parameter `1` is being set correctly. So that's not the problem.

The problem is that *within* a function the positional parameters (from `1` onward, `0` doesn't change) are set to the arguments that the function was given when it was called. (Similarly, within a function the special parameters `@` and `#` are set to all the arguments passed to the function, and the number of arguments passed to the function, respectively.) Since we called the function `hello` without any arguments, while the function `hello` is executing positional parameter `1` is unset, and so when we try to print its value, nothing is printed.

The way you call a shell function with arguments is to list those arguments immediately after the name of the shell function, e.g. in our script:

```
hello Dave
```

would call the function `hello` with one argument: "Dave".

## Shell function arguments (3)

```
> gedit hello-function.sh &
#!/bin/bash
function hello()
{
    # This is a shell function.
    echo "Hello, ${1}"
    echo "I am function ${FUNCNAME}."
}

echo "First argument: ${1}"
hello Hal

> ./hello-function.sh Dave
First argument: Dave
Hello, Hal
I am function hello.
```

[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

30

Modify the file `hello-function.sh` in your home directory as shown above. *Make sure you save the file after you've modified it or your changes won't take effect.*

So, if we call our function with an argument (in this case the argument is "Hal"), then the value of the positional parameter `1` is indeed set to that argument within the function.

So, if we want to our function to have the same first argument as the shell script itself, then we need to call the function with the first argument with which the shell script was invoked.

You can probably guess how we do this...

## Shell function arguments (4)

```
> gedit hello-function.sh &
#!/bin/bash
function hello()
{
    # This is a shell function.
    echo "Hello, ${1}"
    echo "I am function ${FUNCNAME}."
}

#echo "First argument: ${1}"
hello "$1"

> ./hello-function.sh Dave
Hello, Dave
I am function hello.
```

[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

31

Modify the file `hello-function.sh` in your home directory as shown above. *Make sure you save the file after you've modified it or your changes won't take effect.*

Note that now we think we've cracked it, we can get rid of our debugging effort. We could delete that line, but, if we were wrong, we'd only have to put it back in again as we tried to figure it out. So it is easier to just comment it out by putting the inserting a hash character (#) at the start of the line - recall that the shell treats everything after a hash at the start of a line as a comment.

But as you probably guessed - it does indeed work the way we want. Positional parameter `1` holds the first argument that was given on the command line to the shell script, so if we want to pass that argument to the `hello` function, we just put:

```
hello "${1}"
```

in our shell script, and voilà!

# Why use shell functions?

- Allow us to structure our shell script:
  - Functions ⇔ sub-tasks
- Easier to write small parts of the shell script at a time:
  - So we can write the easy bits first!
- Easier to test individual parts of the shell script
- Repetitive sequences of commands only appear in one place:
  - Less typing! Fewer typos!
  - Easy to make changes
  - Easy to fix errors
- Can re-use functions in different shell scripts

[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

32

If you've implemented your shell script entirely as shell functions, there is a really nice trick you can use when something goes wrong and you need to debug your script, or if you want to re-use some of those functions in another script. As you've implemented the script entirely as a series of functions, you have to call one of those functions to start the script actually doing anything. For the purposes of this discussion, let's call that function `main`. So your script looks something like:

```
function start()
{
    ...
}
function do_something()
{
    ...
}
function end()
{
    ...
}
function main()
{
    ...
}

main
```

By commenting out the call to the `main` function, you now have a shell script that does *nothing* except define some functions. You can now easily call the function(s) you want to debug/use from another shell script using the `source` shell builtin command (as we'll see in the "Unix Systems: Shell Scripting (III)" course). This makes debugging *much* easier than it otherwise might be, even of really long and complex scripts.

## Improving `multi-run.sh`

- Depends on `run-once.sh`
- Location of `run-once.sh` hard-coded into script:
  - If we move `run-once.sh`, script breaks until it is updated

```
#!/bin/bash

# Parameters that stay the same each run
myFIXED_PARAMS="100 100 1000"

# Run iterator program once for each argument
# Note: *no* quotes around ${myFIXED_PARAMS}
#       or they'll be interpreted as one argument!
for zzARGS in "${@}" ; do
    ~/scripts/run-once.sh ${myFIXED_PARAMS} ${zzARGS}
done
```

[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

33

In the “Unix Systems: Shell Scripting (I)” course we met the scripts `multi-run.sh` and `run-once.sh`. Together, these scripts gave us a nice way of running a program several times with different parameter sets. However, they are not as versatile as we might hope. `run-once.sh` requires that the program it runs (*iterator*) be in the current directory. Since, in this example, *iterator* is a special program for us (imagine it were your program that you had written from scratch), that’s not such a bad limitation, since we quite probably would have a working copy of the program in the directory where we were going to store its output.

`multi-run.sh`, on the other hand, depends on the `run-once.sh` script, and has the location of that script *hard-coded* into it. If we move the `run-once.sh` script for some reason, then `multi-run.sh` will immediately stop working. Wouldn’t it be nice if we could somehow avoid this problem, but still keep the functionality of the two scripts somewhat separate?

One of way of doing exactly that would be to incorporate `run-once.sh` into `multi-run.sh` as a shell function. That should be quite easy. We define a function in `multi-run.sh` that does *exactly* the same thing as the `run-once.sh` script, and in our `for` loop, instead of calling the `run-once.sh` script, we call our function.

So, let’s do that and see what happens.

# First exercise

Add `run-once.sh` to `multi-run.sh` as a *shell function*:

```
#!/bin/bash

function run_program()
{
# This function runs the iterator program
                        What goes here?
}

# Parameters that stay the same each run
myFIXED_PARAMS="100 100 1000"

# Run iterator program once for each argument
# Note: *no* quotes around ${myFIXED_PARAMS}
#       or they'll be interpreted as one argument!
for zzARGS in "${@"}"; do
    run_program ${myFIXED_PARAMS} ${zzARGS}
done
```

[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

34

The `multi-run.sh` and `run-once.sh` shell scripts are in the `scripts` directory of your home directory. Your task is to get the functionality of the `run-once.sh` script into the `multi-run.sh` script as a *shell function*. Above I've given you the skeleton of what the modified script should look like. You should be able to fill in the rest.

This should be a quick exercise, so when you finish it, take a short break and then we'll start again with the solution. (I really *do* mean take a break – sitting in front of computers for long periods of time is very bad for you. Get up, move around, have a drink, do a little dance, relax...)

You can check that you've done it correctly by trying to run your modified `multi-run.sh` script (*remember to save it after you've made your modifications!*):

```
> cd
> rm -f *.dat stdout-* logfile
> ls
> scripts/multi-run.sh 0.05 0.1 0.15 0.2 0.25 0.3 0.35 0.4 0.45 0.5
> ls
```

*Hint:* This exercise is essentially a cut-and-paste (or copy-and-paste) task. If you are trying to do anything much more complicated than that, then you're on the wrong track...

## Recap: Shell functions

- “mini-shell scripts”
- Usually used for well-defined tasks (often called repeatedly)
- Specify arguments by listing them after function name when calling function  
`hello Dave`
- Positional parameters (and related special shell parameters) set to function’s arguments within function  
In function `hello`, positional parameter 1 = Dave

One thing worth noticing from the exercise we’ve just done:

The original script had the line:

```
~/scripts/run-once.sh ${myFIXED_PARAMS} ${zzARGS}
```

The new script has the line:

```
run_program ${myFIXED_PARAMS} ${zzARGS}
```

Note that arguments that we are passing have not changed in the slightest. In the original script we were calling another shell script with some arguments. In our new script we are calling a *shell function* with the *same* arguments. The syntax for these is almost identical: the main change is the name (and location) of the things being called. See?, I told you shell functions were like “mini-shell scripts”. ☺

# Testing

```
#!/bin/bash

function run_program()
{
# This function runs the iterator program
    ...
}

# Parameters that stay the same each run
myFIXED_PARAMS="100 100 1000"

run_program ${myFIXED_PARAMS} 0.8

# Run iterator program once for each argument
# Note: *no* quotes around ${myFIXED_PARAMS}
#       or they'll be interpreted as one argument!
#for zzARGS in "${@}" ; do
#  run_program ${myFIXED_PARAMS} ${zzARGS}
#done
```

[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

36

One of the advantages of writing a shell script using shell functions should be immediately apparent. The main body of this shell script – the `for` loop – is nice and simple. It just calls a function over and over varying one parameter each time. Because we’ve hidden the commands that do the real work in a shell function, we can see this immediately just by looking at the script.

If we’d put all the lines in the `run_program` function in the `for` loop it would have obscured the script’s structure, and we might have spent a lot of time trying to figure out what the individual lines of script did before realising what was going on. It also helps that we’ve chosen a meaningful name for our shell function. So just by looking at the script we can immediately say “Aha! This script probably runs a program (`run_program`) several times, varying one of its parameters each time.” (Of course, at this point we’d be taking it on faith that the author of the shell script wasn’t an evil troll who deliberately chose *misleading* names for his shell functions. Fortunately, most of those spend the majority of their time under bridges harassing goats.)

Another advantage is that we can easily test our shell function by just commenting the other complicating bits of the shell script out (as above) and just running the function once with some test arguments. This is worth doing every time you’ve written a new function (especially if it is complicated) so that you know it behaves the way you expected it to. It also means that you know that, if there is an error, it is *not* in that part of the shell script (that shell function). That makes it much easier to track down errors.

You can save the above modifications and try out the script if you want: it should just run the `run_program` function once, producing two output files (`output-0.8.dat` and `stdout-0.8`) and writing some information about what it is doing to the log file `logfile`.

If you do try it out, make sure that you undo those modifications and return the shell script to its former state (and save it) as we will be using the shell script later.

# Command substitution

Sometimes we want to get the output of a command and use it in our shell script, for instance, we might want a shell variable to hold the output of a command. How do we do this?:

```
$( command )
```

```
> cd /tmp
> myDIRECTORY="$(pwd)"
> echo "I will use directory: ${myDIRECTORY}"
I will use directory: /tmp
```

[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

37

*Command substitution* is the process whereby the shell runs a *command* and *substitutes* the command's output for wherever the command originally appeared (in a shell script or on the command line).

So, for example, the following line in a shell script:

```
myDIR="$(pwd)"
```

would set the shell variable `myDIR` to the full path of the current working directory. (We don't have to surround the `$(pwd)` in quotes, but it is a good idea: the path may contain spaces.) This is how it works:

1. The shell runs the `pwd` command. The `pwd` command prints out the full path of current working directory, i.e. its output is the full path of the current working directory. Let's suppose we were in `/tmp`, so the output of the `pwd` command would be `/tmp`.
2. The shell takes this output (`/tmp`) and substitutes it for where the original expression `$(pwd)` appeared. So what we now have is:

```
myDIR="/tmp"
```

3. As you probably by now know, this is just the normal way of assigning a value to a shell variable, and, sure enough that's exactly what the shell does: it assigns the value `/tmp` to the shell variable `myDIR`.

Instead of the `$( )` construct you can also use *backquotes*, i.e. you can use ``command`` instead of `$(command)`, and you are likely to come across these in many shell scripts. However, the use of backquotes is generally a very bad idea for two reasons: (1) it's very easy to misplace or overlook a backquote (with catastrophic results) as the backquote character (```) is so small, and (2) it's very difficult to use backquotes to do nested command substitution.

## Improving multi-run.sh (2)

```
#!/bin/bash

function run_program()
{
    ...
    # Run program with passed arguments
    "myPROG" "${@}" > "stdout-${4}"
    ...
}

# Program to run: iterator
myPROG="$(pwd -P)/iterator"

# Parameters that stay the same each run
myFIXED_PARAMS="100 100 1000"
...
```

[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

38

Let's make a small, but major, improvement to the `multi-run.sh` script (this script is in the `scripts` directory of your home directory).

Change the lines:

```
# Run iterator with passed arguments
./iterator "${@}" > "stdout-${4}"
```

to:

```
# Run program with passed arguments
"myPROG" "${@}" > "stdout-${4}"
```

And add the lines:

```
# Program to run: iterator
myPROG="$(pwd -P)/iterator"
```

Immediately *before* the line:

```
# Parameters that stay the same each run
```

Why is this such a major improvement?

Firstly, by replacing the hard-coded `./iterator` with a shell variable, we have made it much easier to modify the script to use other programs instead of `iterator`. (Not to mention making it much more obvious where we make such a modification.)

Secondly, by obtaining the full path of the `iterator` program our shell script can now work in another directory than the one we start off in, as we now have a full path to the `iterator` program and so can run it whatever directory we may be in. We'll see why this is a good idea in a minute.

For those wondering what "`pwd -P`" does, recall that, as already mentioned, `pwd` prints out the full path of the current working directory. With the `-P` option it prints out the full *physical* path, i.e. the path will contain no symbolic links which might change over time (or otherwise confuse things).

You can check that this modified `multi-run.sh` script still works – *remember to save it after you've made your modifications* – with the same sequence of commands given for this purpose on the page 34 of your notes.

# The `mktemp` command

Safely *makes temporary* files or directories for you

## Options:

- d make a *directory* instead of a file
- t make file or directory in a *temporary* directory (usually `/tmp`)

```
> mktemp -t -d iterator.XXXXXXXXXX  
/tmp/iterator.khhcE30735
```

The `mktemp` command is an extremely useful command that allows users to *safely* create temporary files or directories on multi-user systems. It is very easy to *unsafely* create a temporary file or directory to work with from a shell script, and, indeed, if your shell script tries to create its own temporary files or directories using the normal Unix commands then it is almost certainly doing so unsafely. Use the `mktemp` command instead.

Note that if you try the example above you will almost certainly get a directory with a different name created for you.

Note also that `mktemp` has more options than the two listed above, but we won't be using them in this course. Note also that if you use a version of `mktemp` earlier than version 1.3 (or a version derived from BSD, such as that shipped with MacOS X) then you can't use the `-t` option, and will have to specify `/tmp` (or another temporary directory) explicitly, e.g.

```
mktemp -d /tmp/iterator.XXXXXXXXXX
```

How do you use `mktemp`? You give it a “template” which consists of a name with some number of X's appended to it, e.g. `iterator.XXXXXX`. `mktemp` then replaces the X's with random letters and numbers to make the name unique and creates the requested file or directory. It outputs the name of the file or directory it has created.

## Improving multi-run.sh (3)

```
#!/bin/bash -e
...
# My current directory
myDIR="$(pwd -P)"
# Temporary directory for me to work in
myTEMP_DIR="$(mktemp -t -d iterator.XXXXXXXXXX)"

# Change to temporary directory
cd "${myTEMP_DIR}"
# Run program once for each argument
# Note: *no* quotes around ${myFIXED_PARAMS}
#       or they'll be interpreted as one argument!
for zzARGS in "${@}" ; do
    run_program ${myFIXED_PARAMS} ${zzARGS}
done

# Copy files back to my directory
cp -fpr . "${myDIR}"
# Go back to my directory
cd "${myDIR}"
# Clean up
rm -rf "${myTEMP_DIR}"
```

[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

40

Modify the `multi-run.sh` script in the `scripts` directory as shown above.

The improvement we've made here is to now do all our calculations in a temporary directory, and only copy the output files (and log file) to our working directory when we've finished.

(You should understand what all the lines of shell script we've just added are doing – if you don't please ask the course giver or demonstrator to explain.)

Why is this an improvement? Well, if, as in this course, the directory we are working from (our home directory) is actually on a network filesystem, then this can have a major impact on performance, particularly when the network is busy (like when a whole classroom is doing this course). By working in `/tmp`, which is usually a local filesystem (as it is for PWF Linux machines) we no longer have to deal with the network overheads and bottlenecks except right at the very end of the process. This should make things much quicker. It also potentially makes things more reliable as well, as it minimises the opportunity for network problems to mess up our work. (Hurrah!)

One other important thing to note is that we've told our script to abort as soon as it hits an error. That's what adding the `-e` to the `#!/bin/bash` at the start of the file does (you did remember to make that modification, right?). Why do this now? The reason is that our shell script is now doing something dangerous: **it is changing the working directory**. Why is that dangerous? Well, imagine I tried to change to a directory and failed for some reason. Thinking I'm in a different directory than I actually am, I promptly delete everything in it. Oops!

You can check that you've done it correctly by trying to run your modified `multi-run.sh` script (*remember to save it after you've made your modifications!*) – with luck, it will be a little bit faster this time!:

```
> cd
> rm -f *.dat stdout-* logfile
> ls
> scripts/multi-run.sh 0.05 0.1 0.15 0.2 0.25 0.3 0.35 0.4 0.45 0.5
> ls
```

# read

Get input from standard input...

...try to put each word (value) in as many separate variables as are provided...

```
read VAR1 VAR2 VAR3
```

## Options:

`-p` Use the following string as a *prompt* for the user

```
> read -p "What is the answer?:" myANSWER
What is the answer?: 42
> echo "${myANSWER}"
42
```

The `read` shell builtin command takes input from standard input (usually the keyboard) and returns it in the specified shell variable. If you don't specify a shell variable, it will return it in a shell variable called `REPLY`.

The `-p` option gives `read` a string that it displays as a prompt for the user.

You can give `read` more than one shell variable in which to return its input. What happens then is that the first *word* it reads goes into the first shell variable, the second word into the second shell variable and so on.

If there are more words than shell variables, the extra words all are put into the last shell variable.

If there are more shell variables than words, each of the extra variables are set to the empty string.

As far as `read` is concerned a "word" is a sequence of characters that does *not* contain a space, i.e. it considers spaces as the thing that separates one word from another. (The technical term for "thing that separates one thing from another" is "*delimiter*".)

# Using read (1)

```
> cd
> gedit scripts/run-once-using-read.sh &

#!/bin/bash

# Read in parameters from standard input
read -p "Input parameters for iterator: " myNX myNY myN_ITER myEPSILON
...

> cd
> scripts/run-once-using-read.sh
Input parameters for iterator: 10 10 100 0.7
>
```

[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

42

In the `scripts` directory there is a shell script called `run-once-using-read.sh`. Open this up with your favourite editor (or `gedit`) and have a look at it.

The first line (that doesn't start with a `#` character) is a `read` shell builtin command that reads some values from standard input and puts them in some shell variables. (You should be able to work out how the rest of the script has been modified to use these shell variables - if there is anything you don't understand, ask the course giver or demonstrator.)

Let's try this script out and see how it behaves.

## Using read (2)

```
> cd
> rm -f *.dat stdout-* logfile *.png
> scripts/run-once-using-read.sh
Input parameters for iterator: 10 10 100 0.7 garbage
> ls
bin                hello.sh          output-0.7 garbage.dat  stdout-0.7 garbage
examples          iterator          run-iterator.sh        treasure.txt
gnuplot           iterator.gplt     scripts
hello-function.sh logfile           source

> gedit scripts/run-once-using-read.sh &

#!/bin/bash

# Read in parameters from standard input
read -p "Input parameters for iterator: " myNX myNY myN_ITER myEPSILON myJUNK

science-support@ucs.cam.ac.uk          Unix Systems: Shell Scripting (II)          43
```

So, on first try it seemed to do what we'd expect. However, if we give it some input that should be invalid something slightly strange happens. If we give it 5 input parameters instead of 4, instead of complaining, or only using the first 4 parameters, it puts the last two parameters together to form one argument ("0.7 garbage") in the above example and runs the `iterator` program with that (we can see this is what is happening by inspecting the contents of the log file `logfile`). The fact that the `iterator` program doesn't complain, as we can see by inspecting the files it produces, is an indication that the `iterator` program is (yet again) not as well written as we might like (an all too common complaint with software).

(Note that the output of the `ls` command may not exactly match what is shown above - in particular there may be other files or directories show, and the colours may be slightly different.)

However, that fact that our script gives it mangled input to work with is an indication that our script is broken. What is the problem and how can we fix it?

Recall how `read` works: if it reads more words (values) than it was given shell variables, it puts all the extra ones together in the last shell variable. This is what is happening here, and it is undesirable. We can fix this by giving `read` an extra "dummy" shell variable that we never use, but that is simply there to hold any extra junk it may read in.

Modify the `run-once-using-read.sh` shell script in the `scripts` directory as shown above (remember to save it when you've finished).

## Using read (3)

```
> cd
> rm -f *.dat stdout-* logfile
> scripts/run-once-using-read.sh
Input parameters for iterator: 10 10 100 0.7 garbage
> ls
bin      hello-function.sh  iterator.gplt  run-iterator.sh  stdout-0.7
examples hello.sh           logfile       scripts          treasure.txt
gnuplot  iterator          output-0.7.dat source
```

[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

44

Now it works better. If we give it more than 4 input parameters it doesn't mangle the 4th argument that it passes to the `iterator` program.

(Note that the output of the `ls` command may not exactly match what is shown above – in particular there may be other files or directories show, and the colours may be slightly different.)

Now this may seem like a lot of trouble to go to for not much in the way of improvement to our script. After all, the original `run-once.sh` script could perfectly well accept a single set of 4 parameters without all these problems – it just wanted them on the command line rather than from standard input.

So, what's the big deal about standard input? After all, if I have lots of parameter sets to run I'm hardly going to sit there and type them all in one at a time!

Well, how many command line arguments can a shell script have? The answer is quite a few but **not** an unlimited number. In fact, if I have thousands of parameter sets, that's definitely going to be too many for me to pass to my shell script all in one go (or even a small number of goes) on the command line. So, how do we deal with situation?

Hmmmm, maybe if I could put all my thousands of parameter sets into a file, and then could somehow get my shell script to read in that file, one parameter set at a time, that might do it... we need to be able to do a few more things to make that particular idea fly, so let's have a look at some of them now...

# Pipes

A *pipe* takes the  
...*output* of one command...  
...and passes it to another  
command as *input*...

```
command1 | command2
```

Pipes can be combined:

```
command1 | command2 | command3
```

A set of one or more pipes is known as a  
*pipeline*

A *pipe* takes the *output* of one command and feeds it to another command as *input*. We tell the shell to do this using the `|` symbol. So:

```
ls | more
```

takes the the output of the `ls` command and passes it to the `more` command, which displays the output of the `ls` command one screenful at a time. We can combine several pipes by taking the output of the last command of each pipe and passing it to the first command in the next pipe, e.g.

```
ls | grep 'fred' | more
```

takes the output of `ls` and passes it to `grep`, which searches for lines with the string “fred” in them, and then the output of `grep` is passed to the `more` command to display one screenful at a time. A set of one or more pipes is known as a *pipeline*. This pipeline would show us all the files with the string “fred” in their name, one screenful at a time.

# Using pipes

```
> cd
> rm -f *.dat stdout-* logfile
> cat scripts/param_set
100 100 1000 0.05
100 100 1000 0.1
100 100 1000 0.15
100 100 1000 0.2
100 100 1000 0.25
100 100 1000 0.3
100 100 1000 0.35
100 100 1000 0.4
100 100 1000 0.45
100 100 1000 0.5
> cat scripts/param_set | scripts/run-once-using-read.sh
> ls
bin      hello-function.sh  iterator.gplt  run-iterator.sh  stdout-0.05
examples hello.sh           logfile        scripts           treasure.txt
gnuplot  iterator          output-0.05.dat source
```

[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

46

In the `scripts` directory there is a file called `param_set` that contains a number of parameter sets. We can use the `cat` command to display the contents of this file. In fact, if we use the `cat` command on this file, the *output* of the `cat` command will be a list of parameter sets...

...and our `run-once-using-read.sh` shell script will accept a complete parameter set as its *input*, so...

...if we connect the *output* of the `cat` command to the *input* of our shell script – by, say, using a pipe – maybe that will give us what we want? Let's try it!

Well, it almost does!, i.e. it does it for the first parameter set, but none of the others. If we try running it again and again it will still only do it for the first parameter set in the file, so we're not quite there, but close. What we want is some way of telling the script to keep reading until there is no more stuff to read.

In fact, what we want is for the script to do some sort of *loop*: reading in a set of values, then running the `iterator` program, then reading in the next set of values, and so on. How can we get it to do that? Before we look at that, we need to understand something else first...

# Exit Status (1)

- Every program (or shell builtin command) returns an *exit status* when it completes
- Number between 0 and 255
- **Not** the same as the program's (or shell builtin command's) output
- By convention:
  - 0 means the command succeeded
  - Non-zero value means the command failed
- Exit status of the last command ran stored in special shell parameter named ?

The exit status of a program is also called its *exit code*, *return code*, *return status*, *error code*, *error status*, *errorlevel* or *error level*.

## Exit Status (2)

```
> ls
bin      hello-function.sh  iterator.gplt  run-iterator.sh  stdout-0.05
examples hello.sh           logfile       scripts          treasure.txt
gnuplot  iterator          output-0.05.dat source

> echo "${?}"
0

> ls zzzfred
/bin/ls: zzzfred: No such file or directory

> echo "${?}"
2
```

You get the value of the special parameter `?` by using the construct `${?}`, as in the above example.

Note that when the `ls` command is successful, its exit status is 0. When, however, it fails (for example because the file does not exist, as here), its exit status is non-zero ("2", in this case). In our shell scripts, we will make significant use of the fact that a non-zero exit status of a program (or a shell builtin command) means that there was an error.

Please note that the output of the `ls` command may not exactly match what is shown on this slide – in particular, the colours may be slightly different shades and there may be additional files and/or directories shown (and/or – if you've recently cleaned up your home directory – you may not have all of the files shown here).

# while

Repeat *while* some expression is *true*

```
while <expression> ; do
    <some commands>
done
```

We can repeat a collection of commands using a `while` loop. A `while` loop repeats a collection of commands as long as the result of some *test* or command is true (what's a test? - we'll meet them in a little while). The result of a command is considered to be true if it returns an exit status of 0 (i.e. if the command succeeded). We use a `while` loop like this:

```
while <expression> ; do
    <some commands>
done
```

where `<expression>` is either a test or a command, and `<some commands>` is a collection of one or more commands.

As with a `for` loop, you can put the `do` on a separate line, in which case you can omit the semi-colon (;).

There are some examples of how to use `while` loops in the following files in the `examples` directory:

```
while1.sh
while2.sh
```

...but don't look at those files just yet as we need to meet a few more things first...

# Using while (1)

```
> cd
> cp -p scripts/run-once-using-read.sh scripts/run-while-read.sh
> gedit scripts/run-while-read.sh &
#!/bin/bash

# Read in parameters from standard input
#   and then run iterator with them
#   and run it again and again until there are no more
while read myNX myNY myN_ITER myEPSILON myJUNK ; do
    ...
echo "Standard output: stdout-${myEPSILON}" >> logfile
done
```

[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

50

Create a copy of the `run-once-using-read.sh` shell script in the `scripts` directory called `run-while-read.sh`. Open this up with your favourite editor (or `gedit`) and modify it as shown above.

Basically, replace the line:

```
read -p "Input parameters for iterator: " myNX myNY myN_ITER myEPSILON myJUNK
```

with:

```
#   and then run iterator with them
#   and run it again and again until there are no more
while read myNX myNY myN_ITER myEPSILON myJUNK ; do
```

And at the very end of the file add the following line:

```
done
```

Remember to save the script when you've finished.

Now let's try this script out and see if it does what we want:

```
> cd
> rm -f *.dat stdout-* logfile
> cat scripts/param_set | scripts/run-while-read.sh
> ls
```

## Second exercise

Make a copy of `multi-run.sh` and make it read ***all*** the arguments for `iterator` in from standard input using a `while` loop:

```
> cd
> cp -p scripts/multi-run.sh scripts/multi-run-while.sh

#!/bin/bash

...

# Run iterator program once for each argument
# Note: *no* quotes around ${myFIXED_PARAMS}
#       or they'll be interpreted as one argument!
for zzARGS in "${@}" ; do
    run_program ${myFIXED_PARAMS} ${zzARGS}
done
```

[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

51

The `multi-run.sh` shell script is in the `scripts` directory of your home directory. Make a copy of it called `multi-run-while.sh`, also in the `scripts` directory, and work on that. Your task is to get `multi-run-while.sh` to read in all the arguments for `iterator` from standard input (all its arguments, not just the fourth one) using a `while` loop.

Start by deleting the following two lines:

```
# Parameters that stay the same each run
myFIXED_PARAMS="100 100 1000"
```

...and you should also get rid of any other references to the shell variable `myFIXED_PARAMS` - you won't be using it in this script.

We have gone through everything you need to do this exercise. You should comment the modifications you make to your shell script, preferably as you are writing it.

And when you finish this exercise, please do take a short break before we start again with the solution. (And that's "break" as in "break from the computer" not "break to check my e-mail".)

You can check that you've done it correctly by trying to run your `multi-run-while.sh` script (*remember to save it after you've made your modifications!*):

```
> cd
> rm -f *.dat stdout-* logfile
> ls
> cat scripts/param_set | scripts/multi-run-while.sh
> ls
```

*Hint:* Try copying the `run-while-read.sh` script...

## Recap: standard input/`while` loops

- Command substitution `$(command)` can be used to get the output of a command into a shell variable
- Use `mktemp` to make temporary files and directories
- `read` gets values from standard input
- Pipes connect one command's output to another's input
- `while` loops repeat some commands while something is true – can be used to read in multiple lines of input with `read`

# Tests

Test to see if something is true:

```
[ <expression> ]
```

or: `test <expression>`

where `<expression>` can be any of a number of things such as:

```
[ "a" -eq "b" ]
```

```
[ "a" -le "b" ]
```

```
[ "a" -gt "b" ]
```

A test is basically the way in which the shell evaluates an expression to see if it is true. (Recall that they can be used with `while` and we'll see how in a minute.)

There are many different tests that you can do, and we only list a few here:

<code>"a" -lt "b"</code>	true if and only if the integer <code>a</code> is less than the integer <code>b</code>
<code>"a" -le "b"</code>	true if and only if the integer <code>a</code> is less than or equal to the integer <code>b</code>
<code>"a" -eq "b"</code>	true if and only if the integer <code>a</code> is equal to the integer <code>b</code>
<code>"a" -ne "b"</code>	true if and only if the integer <code>a</code> is not equal to the integer <code>b</code>
<code>"a" -ge "b"</code>	true if and only if the integer <code>a</code> is greater than or equal to the integer <code>b</code>
<code>"a" -gt "b"</code>	true if and only if the integer <code>a</code> is greater than the integer <code>b</code>

You can often omit the quotation marks, particularly for arithmetic tests (we'll meet other sorts of tests in "Unix Systems: Shell Scripting (III)"), but it is good practice to get into the habit of using them, since there are times when *not* using them can be disastrous.

In the above tests, `a` and `b` can be any integers. Recall that shell variables can hold pretty much any value we like – they can certainly hold integer values, so `a` and/or `b` in the above expressions could come from shell variables, e.g.

```
[ "${VAR}" -eq "5" ]
```

Or, equivalently:

```
test "${VAR}" -eq "5"
```

is true if and only if the shell variable `VAR` contains the value "5".

## Arithmetic Expansion: `$( ( ) )`

- Returns the value of an *integer* arithmetic operation
- Operands *must* be integers (so *no* decimals, e.g. 2.5, etc)
- Do *not* use quotes in the arithmetic expression

```
$( ( <arithmetic-expression> ) )
```

Example:

```
$( ( ${VAR} + 56 ) )
```

The shell can also do (primitive) integer arithmetic, which can be very useful, as we will see in a minute.

The construct `$( (<arithmetic-expression> ) )` means replace `$( (<arithmetic-expression> ) )` with the result of the *integer* arithmetic expression `<arithmetic-expression>`. This is known as *arithmetic expansion*. (The arithmetic expression is evaluated as integer arithmetic.) Note that we *don't* use quotes around our variables in our arithmetic expression as that would cause the shell to treat the values as strings rather than numbers (this is, alas, somewhat inconsistent with the shell's behaviour elsewhere).

## while loops that count

Consider the following `while` loop:

```
zzCOUNT=0
while [ "${zzCOUNT}" -le "5" ] ; do
    zzCOUNT=$(( ${zzCOUNT} + 1 ))
    echo "${zzCOUNT}"
done
```

When we put together arithmetic tests, `while` loops and arithmetic expansion, we can construct a `while` loop that counts for us, as in the above example. Can you figure out what the above loop will do?

When you think you know, try running the script `while2.sh` in the `examples` directory of your home directory. That will show you the output of the above `while` loop, immediately followed by the output of a very similar `while` loop where `zzCOUNT` starts off with the value 1 rather than 0.

## Using while (2)

```
> cd
> gedit scripts/generate-params.sh &

#!/bin/bash

myNX="10"
myNY="10"
myEPSILONS="0.1 0.2 0.3 0.4 0.5"

for zzEPS in ${myEPSILONS} ; do
    zzITERS="10"
    while [ "${zzITERS}" -le "10000" ] ; do
        echo "${myNX} ${myNY} ${zzITERS} ${zzEPS}" >> new_param_set
        zzITERS=$(( ${zzITERS} * 10 ))
    done
done

> cd
> chmod +x scripts/generate-params.sh
> scripts/generate-params.sh
> more new_param_set
```

[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

56

Create a new file called `generate-params.sh` in the `scripts` directory of your home directory and type the shell script shown above into it.

Remember to save the script when you've finished and to set the execute bit on it using the `chmod` command.

Then try it out and see what it does.

## Generalising multi-run-while.sh

```
#!/bin/bash -e

...

"${myPROG}" "${@"}" > "stdout--${1}-${2}-${3}-${4}"

...

mv output.dat "output--${1}-${2}-${3}-${4}.dat"

...

echo "Output file: output--${1}-${2}-${3}-${4}.dat" >> logfile
echo "Standard output: stdout--${1}-${2}-${3}-${4}" >> logfile
```

Modify the `multi-run-while.sh` script in the `scripts` directory as shown above.

(Remember to save it when you've finished.)

Basically we are replacing all the instances of the string “`-${4}`” with the string “`-${1}-${2}-${3}-${4}`”. This means that now, instead of our output files being based on the fourth argument that is passed to `iterator`, they are based on all the parameters in the parameter set. This is clearly necessary as we start to experiment with varying parameters other than just the fourth one.

And we finish with an exercise.

If you want to do the exercise outside of class, the files you'll need can be found at:

<http://www-uxsup.csx.cam.ac.uk/courses/ShellScriptingII/exercises.html>

## Final exercise – Part One

Improve the `run_program` function in `multi-run-while.sh` so that as well as running `iterator` it also runs `gnuplot` (using the `iterator.gplt` file) to plot a graph of the output.

This exercise should be fairly straightforward. One sensible way of approaching it would be as follows:

1. Figure out the full path of the `iterator.gplt` file. Store it a shell variable (maybe called something like `myGPLT_FILE`).
2. Immediately after running `iterator`, run `gnuplot`:

```
gnuplot "${myGPLT_FILE}"
```
3. Rename the `output.png` file produced by `gnuplot` along the same lines as the `output.dat` file produced by `iterator` is renamed.

***Make sure you test the script after you've modified it and check that it does what you would expect.***

This exercise highlights one of the advantages of using functions: we can improve or change our functions whilst leaving the rest of the script unchanged. In particular, the *structure* of the script remains unchanged. This means two things: (1) if there are any errors after changing the script they are almost certainly in the function we changed, and (2) the script is still doing the same *kind* of thing (as we can see at a glance) – we've just changed the particulars of one of its functions.

## Final exercise – Part Two

Now create a new shell script based on `multi-run-while.sh` that will run `iterator` three times **for** each **parameter set** the script reads in on standard input, **changing** the third parameter each time as follows:

For a given parameter set `a b c d`, first your script should run `iterator` with the parameter set:

`a b 10 d`

...then with the parameter set:

`a b 100 d`

...and then with the parameter set:

`a b 1000 d`

An example may help to make this task clearer. Suppose your script reads in the parameter set:

`10 10 50 0.5`

...it should then run the `iterator` program 3 times, once for each of the following parameter sets:

`10 10 10 0.5`

`10 10 100 0.5`

`10 10 1000 0.5`

The first thing to do is to make a copy of the `multi-run-while.sh` script and work on the copy – I suggest you call your copy something like `multi-10-100-1000.sh`:

```
> cd
```

```
> cp -p scripts/multi-run-while.sh scripts/multi-10-100-1000.sh
```

Now, currently the script will read in a parameter set and then call the `run_program` function to process that parameter set. Clearly, instead of passing all four parameters that the script reads in, your new script will now only be passing the first (`myNX`), second (`myNY`), and fourth (`myEPSILON`) parameters that it has read in. However, the `iterator` program *requires* 4 parameters (and it cares about the order in which you give them to it), so your script still needs to give it 4 parameters, it is just going to ignore the third parameter it has read (`myN_ITER`) and substitute values of its own instead.

There are two approaches you could take. One would be to call the `run_program` function 3 times, once with 10 as the third parameter, once with 100 as the third parameter and once with 1000 as the third parameter. The other would be to use some sort of loop that calls the `run_program` function, using the appropriate value (10, 100 or 1000) for the third parameter on each pass of the loop. **I want you to use the loop approach.**

*Hint:* Use a `for` loop.

## Final exercise – Part Three

Now create a new shell script, based on the script you created in the previous part of the exercise, that does the following:

Instead of running `iterator` three times for each parameter set it `reads` in, this script should accept a set of values on the command line, and use those instead of the hard-coded 10, 100, 1000 previously used.

Thus, for each parameter set it `reads` in on standard input, it should run `iterator` substituting, in turn, the values from the command line **for** the third parameter in the parameter set it has `read` in.

So, if the script from the previous part of the exercise was called `multi-10-100-1000.sh`, and we called this new script `multi-iterations.sh` (and stored both in the `scripts` directory of our home directory), then running the new script like this:

```
> cat ~/scripts/param_set | ~/scripts/multi-iterations.sh 10 100 1000
    should produce exactly the same output as running the old script
    with the same input file:
> cat ~/scripts/param_set | ~/scripts/multi-10-100-1000.sh
```

[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

60

The first thing to do is to make a copy of the previous script (which I suggested you call `multi-10-100-1000.sh`) and work on the copy – I suggest you call your copy something like `multi-iterations.sh`:

```
> cd
> cp -p scripts/multi-10-100-1000.sh scripts/multi-iterations.sh
```

You may be wondering what the point of the previous script and this script are. Consider what these scripts actually do: they take a parameter set, vary one of its parameters and then run some program with the modified parameter sets. Why would we want to do this?

Well, in this example the parameter we are varying specifies the number of iterations for which our program will run. You can easily imagine that we might have a simulation or calculation for which, for any given parameter set, interesting things happened after various numbers of iterations. These scripts allow us to take each parameter set and run it several times for different numbers of iterations. We can then look at each parameter set and see how varying the number of iterations affects the program's output for that parameter set.

If we were using the parameter sets in the `scripts/param_set` file, we might notice that these parameters are the same except for the fourth parameter which varies. So if we pipe those parameter sets into one of these scripts, we are now investigating how the output of the `iterator` program varies as we vary *two* of its input parameters, which is kinda neat, doncha think? ☺

*Hint:* Modify the loop you used in the previous script to loop over all the command line arguments rather than some hard coded values. If you don't remember the construct that gives you all the command line arguments have a look at the recap of the "Shell Scripting (I)" course.

## Final exercise – Files

All the files (scripts, `iterator` program etc) used in this course are available on-line at:

<http://www-uxsup.csx.cam.ac.uk/courses/ShellScriptingII/exercises.html>

We'll be looking at the answers to this final exercise in the “Unix Systems: Shell Scripting (III)” course, so if you are attending that course (and I recommend that you do if you can) then *please make sure you have attempted this exercise **before** you come to that course.*

Details of the “Unix Systems: Shell Scripting (III)” course can be found on-line at:

<http://www.cam.ac.uk/cs/courses/coursedescriptwkschp>

## Appendix: Unix commands (1)

`cat`            Display contents of a file

```
> cat /etc/motd
```

```
Welcome to PWF Linux 2006/2007.
```

```
If you have any problems, please email Help-Desk@ucs.cam.ac.uk.
```

`cd`            **change directory**

```
> cd /tmp
```

```
> cd
```

`chmod`        **change the mode** (permissions) of  
a file or directory

```
> chmod a+r treasure.txt
```

[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

62

If you give the `cd` command without specifying a directory then it will change the directory to your *home directory* (the location of this directory is specified in the `HOME` environment variable).

The `chmod` command changes the permissions of a file or directory (in this context, the jargon word for “permissions” is “mode”). For instance, the above example gives read access to the file `treasure.txt` for all users on the system. Unix permissions were covered in the “Unix System: Introduction” course, see:

<http://www.cam.ac.uk/cs/courses/coursedescript/linux.html#unix>

## Appendix: Unix commands (2)

`cp` **copy** files and/or directories

```
> cp /etc/motd /tmp/motd-copy
```

### Options:

- p **p**reserve (if possible) files' owner, permissions & date
- f if unable to overwrite destination file, delete it and try again, i.e. **f**orcibly overwrite destination files
- r copy any directories **r**ecursively, i.e. copy their contents

```
> cp -p /etc/motd /tmp/motd-copy
```

Note that the `cp` command has many other options than the three listed above, but those are the options that will be most useful to us in this course.

## Appendix: Unix commands (3)

`date`          display/set system *date* and time

```
> date
```

```
Fri Feb 16 11:52:03 GMT 2007
```

`echo`          display text

```
> echo "Hello"
```

```
Hello
```

`env`            With no arguments, display  
*environment* variables

Please note that if you try out the `date` command, you will get a different date and time to that shown on this slide (unless your computer's clock is wrong or you have fallen into a worm-hole in the space-time continuum). Also, note that usually only the system administrator can use `date` to set the system date and time.

Note that the `echo` command has a few useful options, but we won't be making use of them today, so they aren't listed.

Note also that the `env` command is a very powerful command, but we will not have occasion to use for anything other than displaying environment variables, so we don't discuss its other uses.

## Appendix: Unix commands (4)

`grep` find lines in a file that match a given pattern

```
> grep 'PWF' /etc/motd
```

```
Welcome to PWF Linux 2006/2007.
```

`ln` create a *link* between files (almost always used with the `-s` option for creating *symbolic* links)

```
> ln -s /etc/motd /tmp/motd
```

```
> cat /etc/motd
```

```
Welcome to PWF Linux 2006/2007.
```

```
If you have any problems, please email Help-Desk@ucs.cam.ac.uk.
```

```
> cat /tmp/motd
```

```
Welcome to PWF Linux 2006/2007.
```

```
If you have any problems, please email Help-Desk@ucs.cam.ac.uk.
```

[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

65

The patterns that the `grep` command uses to find text in files are called *regular expressions*. We won't be covering these in this course, but if you are interested, or if you need to find particular pieces of text amongst a collection of text, then you may wish to attend the CS "Pattern Matching Using Regular Expressions" course, details of which are given here:

<http://www.cam.ac.uk/cs/courses/coursedescript/linux.html#regex>

The `ln` command creates links between files. In the example above, we create a symbolic link to the file `motd` in `/etc` and then use `cat` to display both the original file and the symbolic link we've created. We see that they are identical.

There are two sort of links: *symbolic links* (also called *soft links* or *symlinks*) and *hard links*. A symbolic link is similar to a shortcut in the Microsoft Windows operating system (if you are familiar with those) – essentially, a symbolic link points to another file elsewhere on the system. When you try and access the contents of a symbolic link, you actually get the contents of the file to which that symbolic link points. Whereas a symbolic link points to another *file* on the system, a hard link points to *actual data* held on the filesystem. These days almost no one uses `ln` to create hard links, and on many systems this can only be done by the system administrator. If you want a more detailed explanation of symbolic links and hard links, see the following Wikipedia articles:

[http://en.wikipedia.org/wiki/Symbolic\\_link](http://en.wikipedia.org/wiki/Symbolic_link)

[http://en.wikipedia.org/wiki/Hard\\_link](http://en.wikipedia.org/wiki/Hard_link)

## Appendix: Unix commands (5)

`ls` *list* the contents of a directory

> **ls**

```
bin examples gnuplot hello.sh iterator scripts source treasure.txt
```

### Options:

- d List *d*irectory name instead of its contents
- l use a *l*ong listing that gives lots of information about each directory entry
- R list subdirectories *R*ecursively, i.e. list their contents and the contents of any subdirectories within them, etc

If you try out the `ls` command, please note that its output may not exactly match what is shown on this slide – in particular, the colours may be slightly different shades and there may be additional files and/or directories shown. Note also that the `ls` command has many, many more options than the three given on this slide, but these three are the options that will be of most use to us in this course.

## Appendix: Unix commands (6)

`less`          Display a file one screenful of text at a time  
`more`          Display a file one screenful of text at a time

**> more treasure.txt**

The Project Gutenberg eBook of Treasure Island, by Robert Louis Stevenson

This eBook is for the use of anyone anywhere at no cost and with almost no restrictions whatsoever. You may copy it, give it away or re-use it under the terms of the Project Gutenberg License included with this eBook or online at [www.gutenberg.org](http://www.gutenberg.org)

Title: Treasure Island

Author: Robert Louis Stevenson

Release Date: February 25, 2006 [EBook #120]

Language: English

Character set encoding: ASCII

\*\*\* START OF THIS PROJECT GUTENBERG EBOOK TREASURE ISLAND \*\*\*

--More-- (0%)

[escience-support@ucs.cam.ac.uk](mailto:escience-support@ucs.cam.ac.uk)

Unix Systems: Shell Scripting (II)

67

*(Note that the output of the `more` command may not exactly match that shown on this slide – in particular, the number of lines displayed before the “--More-- (0%)” message depends on the number of lines it takes to fill up the window in which you are running the `more` command.)*

The `more` and `less` commands basically do the same thing: display a file one screenful of text at a time. Indeed, on some Linux systems the `more` command is actually just another name (an *alias*) for the `less` command.

Why are there two commands that do the same thing? On the original Unix systems, the `less` command didn't exist – the command to display a file one screenful of text at a time was `more`. However, the original `more` command was somewhat limited, so someone wrote a better version and called it `less`. These days the `more` command is a bit more sophisticated, although the `less` command is still much more powerful.

For everyday usage though, many users find the two commands are equivalent. Use whichever one you feel most comfortable with, but remember that every Unix/Linux system should have the `more` command, whereas some (especially older Unix systems) may not have the `less` command.

## Appendix: Unix commands (7)

`mkdir` *make directories*

```
> mkdir /tmp/mydir
```

Options:

`-p` make any *p*arent directories as required;  
also if directory already exists, don't  
consider this an error

```
> mkdir /tmp/mydir
```

```
mkdir: cannot create directory `/tmp/mydir': File exists
```

```
> mkdir -p /tmp/mydir
```

Note that the `mkdir` command has other options, but we won't be using them in this course.

## Appendix: Unix commands (8)

`mv` **move** or rename files and directories

```
> mv /tmp/motd-copy /tmp/junk
```

### Options:

- f do not prompt before overwriting files or directories, i.e. **f**orcibly move or rename the file or directory; this is the default behaviour
- i prompt before overwriting files or directories (be **i**nteractive - ask the user)
- v show what is being done (be **v**erbose)

Note that the `mv` command has other options, but we won't be using them in this course. Note also that if you move a file or directory between different filesystems, `mv` actually copies the file or directory to the other filesystem and then deletes the original.

## Appendix: Unix commands (9)

`pwd` **p**rint full path of current **w**orking  
**d**irectory

```
> cd /tmp
```

```
> pwd
```

```
/tmp
```

### Options:

- P print the full **P**hysical path of the current working directory (i.e. the path printed will not contain any symbolic links)

Note that the `pwd` command has another option, but we won't be using it in this course.

## Appendix: Unix commands (10)

`rm` **re**move files or directories

> `rm /tmp/junk`

### Options:

- f ignore non-existent files and do not ever prompt before removing files or directories, i.e. **f**orcibly remove the file or directory
- i prompt before removing files or directories (be **i**nteractive - ask the user)
- preserve-root do not act recursively on /
- r remove subdirectories (if any) **r**ecursively, i.e. remove subdirectories and their contents
- v show what is being done (be **v**erbose)

Note that the `rm` command has other options, but we won't be using them in this course.

## Appendix: Unix commands (11)

`rmdir`    *re*move *empty* *dir*ectories

> **`rmdir /tmp/mydir`**

`touch`    change the timestamp of a file;  
if the file doesn't exist create it  
with the specified timestamp  
(the default timestamp is the  
current date and time)

> **`touch /tmp/nowfile`**

The `rmdir` and `touch` commands have various options but we won't be using them on this course. If you try out the `touch` command with the example above, check that it has really worked the way we've described here by using the `ls` command as follows:

```
ls -l /tmp/nowfile
```

You should see that the file `nowfile` has a timestamp of the current time and date.