



Image Sharpening

Practical Introduction to HPC Exercise

(ARCHER version)



1 Aims

The aim of this exercise is to get you used to logging into an HPC resource, using the command line and an editor to manipulate files, and using the batch submission system. We will be using ARCHER for this exercise. ARCHER is the UK national HPC service, and is a Cray XC30 system with a total of 118,010 cores (4920 nodes).

You can find more details on ARCHER and how to use it in the User Guide at:

- <http://www.archer.ac.uk/documentation/user-guide/>

2 Introduction

In this exercise you will run a simple program, in serial and parallel, to sharpen the provided image. Using your provided guest account, you will:

1. log onto the ARCHER frontend nodes;
2. copy the source code from a central location to your account;
3. unpack the source code archive;
4. compile the source code to produce an executable file;
5. run a serial job on the login node;
6. submit a serial job to the compute nodes using the PBS batch system;
7. submit a parallel job using the PBS batch system;
8. run the parallel executable using an increasing number of cores and examine the performance improvement.

Please **do ask questions** in the tutorials if you do not understand anything in the instructions.

3 Instructions

3.1 Log into ARCHER frontend nodes and run commands

You should use your ARCHER guest user name and password to log into ARCHER.

3.1.1 Procedure for Mac and Linux users

Open a command line *Terminal* and enter the following command:

```
ssh -XY user@login.archer.ac.uk
```

You should be prompted to enter your password.

3.1.2 Procedure for Windows users

Windows does not generally have SSH installed by default, so some extra work is required. We recommend you install the free MobaXterm application <https://mobaxterm.mobatek.net/> which incorporates an SSH client, an X11 server for displaying Unix graphics and a drag-and-drop file browser. If you do not have the ability to install software on your machine (e.g. it is a managed desktop) then you should install the Portable edition.

After starting MobaXterm you should click on Sessions -> New Session from the top bar and then select "SSH". You should enter the ARCHER login address (login.archer.ac.uk) as the "Remote host".

When you connect you will be prompted for your username and password.

3.2 Running commands

You can list the directories and files available by using the *ls* (LiSt) command:

```
input
ls
output
bin work
```

NB: The first time you do this there will be no files or directories so 'ls' will return with an empty line.

You can modify the behaviour of commands by adding options. Options are usually letters or words preceded by '-'. For example, to see more details of the files and directories available you can add the '-l' (l for long) option to *ls*:

```
input
ls -l
output
total 8
drwxr-sr-x 2 user z01 4096 Nov 13 14:47 bin
drwxr-sr-x 2 user z01 4096 Nov 13 14:47 work
```

If you want a description of a particular command and the options available you can access this using the *man* (MANual) command. For example, to show more information on *ls*:

```
input
man ls
output
Man: find all matching manual pages
* ls (1)
ls (lp)
Man: What manual page do you want?
Man:
```

In the manual, use the spacebar to move down, 'u' to move up, and 'q' to quit to the command line.

3.3 Download and extract the exercise files

Firstly, change directory so you are on the “/work” file system on ARCHER.

```
input
cd /work/y14/y14/user/
```

You should replace “user” by your user name, e.g. guest23. /work is a high performance parallel file system that can be accessed by both the frontend and compute nodes. All jobs on ARCHER should be run from the /work file system. ARCHER compute nodes cannot access the /home file system at all. Any jobs attempting to use /home will fail with an error.

Use *wget* (on ARCHER) to get the exercise files archive from the EPCC webserver. Material for the course is stored at:

<http://www.archer.ac.uk/training/course-material/2019/07/intro-epcc/>

If you access the material with a web browser it will be downloaded to your laptop, which is not very useful if we want the files on ARCHER. It is easier to copy them straight from the website using the *wget* utility on ARCHER:

Input

```
wget http://www.archer.ac.uk/training/course-material/2019/07/intro-epcc/exercises/sharpen.tar.gz
```

output

```
--2014-06-27 16:15:42--
http://www.archer.ac.uk/training/course-material/...
Resolving www.archer.ac.uk... 193.62.216.12
Connecting to www.archer.ac.uk|193.62.216.12|:80... connected.
HTTP request sent, awaiting response... 200 OK
Length: 1754173 (1.7M) [application/x-gzip]
Saving to: 'sharpen.tar.gz'
100%[=====>] 1,754,173  --.-
K/s in 0.02s
```

To unpack the archive:

input

```
tar -xzvf sharpen.tar.gz
```

output

```
sharpen/C-SER/
sharpen/C-SER/filter.c
...
sharpen/F-OMP/dosharpen.f90
sharpen/F-OMP/Makefile
sharpen/F-OMP/fuzzy.pgm
```

If you are interested in the C examples move to the *C-SER* subdirectory; for Fortran, move to *F-SER*. For example:

```
cd sharpen/C-SER
```


3.4 Using the Emacs text editor

Running interactive graphical applications on ARCHER can be very slow. It is therefore best to use Emacs in *in-terminal* mode. In this mode you can edit the file as usual but you must use keyboard shortcuts to run operations such as “save file” (remember, in this mode there are no menus that can be accessed using the mouse).

Start Emacs with the `emacs -nw` command and the name of the file you wish to edit or create

```
emacs -nw sharpen.pbs
```

The terminal will change to show that you are now inside the Emacs text editor:



Typing will insert text as you would expect and backspace will delete text. You use special key sequences (involving the Ctrl and Alt buttons) to save files, exit Emacs and so on.

Files can be saved using the sequence “Ctrl-x Ctrl-s” (usually abbreviated in Emacs documentation to “C-x C-s”). You should see the following briefly appear in the line at the bottom of the window (the minibuffer in Emacs-speak):

```
Wrote ./sharpen.pbs
```

To exit Emacs and return to the command line use the sequence “C-x C-c”. If you have changes in the file that have not yet been saved Emacs will prompt you (in the minibuffer) to ask if you want to save the changes or not.

Although you could edit files on your local machine using whichever windowed text editor you prefer it is useful to know enough to use an in-terminal editor as there will be times where you want to perform a quick edit that does not justify the hassle of editing and re-uploading.

3.5 Other text editors

Although emacs is a very widely used and powerful text editor, it is not particularly easy for new users. On ARCHER, you should also be able to use “gedit” which is a simple graphical editor.

MobaXterm has a built-in editor which you can launch by clicking on files in the file browser. Note that to access your work directory from the file browser you will need to enter “/work/y14/y14/user/” directly into the location bar.

3.6 Useful commands for examining files

There are a couple of commands that are useful for displaying the contents of plain text files on the command line that you can use to examine the contents of a file without having to open in Emacs (if you want to edit a file then you will need to use Emacs). The commands are *cat* and *less*. *cat* simply prints the contents of the file to the terminal window and returns to the command line. For example:

```
cat sharpen.pbs
...
```

This is fine for small files where the text fits in a single terminal window. For longer files you can use the *less* command: *less* gives you the ability to scroll up and down in the specified file. For example:

```
less Makefile
```

Once in *less* you can use the spacebar to scroll down and ‘u’ to scroll up. When you have finished examining the file you can use ‘q’ to exit *less* and return to the command line.

This program takes a fuzzy image and uses a simple algorithm to sharpen the image. A very basic parallel version of the algorithm has been implemented which we will use in this exercise. There are a number of versions of the sharpen program available:

- C-SER Serial C version
- F-SER Serial Fortran version
- C-MPI Parallel C version using MPI
- F-MPI Parallel Fortran version using MPI
- C-OMP Parallel C version using OpenMP
- F-OMP Parallel Fortran version using OpenMP

3.7 Compile the source code to produce an executable file

We will first compile the serial version (using the *make* command) of the code for our example.

input

```
ls
```

output

```
cio.c          filter.c  Makefile  sharpen.h  utilities.c
dosharpen.c   fuzzy.pgm sharpen.c  sharpen.pbs utilities.h
```

input

```
make
```

output

```
cc -g -DC_SERIAL_PRACTICAL -c sharpen.c
cc -g -DC_SERIAL_PRACTICAL -c dosharpen.c
cc -g -DC_SERIAL_PRACTICAL -c filter.c
cc -g -DC_SERIAL_PRACTICAL -c cio.c
cc -g -DC_SERIAL_PRACTICAL -c utilities.c
cc -g -DC_SERIAL_PRACTICAL -o sharpen sharpen.o dosharpen.o
filter.o cio.o utilities.o
```

This should produce an executable file called *sharpen* which we will run on ARCHER. For the Fortran version, the process is exactly the same as above, except you should move to the *F-SER* subdirectory and build the program there:

input

```
cd sharpen/F-SER
make
```

As before, this should produce a *sharpen* executable. Don't worry about the C file *utilities.c* – it is just providing an easy method for printing out various information about the program at run time, and it is most easily implemented in C rather than Fortran.

3.8 Running a serial job

You can run this serial program directly on the login nodes, e.g.:

input

```
./sharpen
```

output

```
Image sharpening code running in serial
Input file is: fuzzy.pgm
Image size is 564 x 770
Using a filter of size 17 x 17
Reading image file: fuzzy.pgm
... done
Starting calculation ...
On core 0-31
... finished
Writing output file: sharpened.pgm
... done
Calculation time was 5.579000 seconds
Overall run time was 5.671895 seconds
```

3.9 Viewing the images

To see the effect of the sharpening algorithm, you can view the images using the display program from the ImageMagick suite.

```
input
    display fuzzy.pgm
    display sharpened.pgm
```

Type “q” in the image window to close the program.

To view the image you will need an X window client installed. This will be available automatically on Linux systems or if you are using MobaXterm on Windows, but Mac users will have to install “Xquartz”.

3.10 Running on the compute nodes

Remember you must be in the “/work” file system on ARCHER – the default is /home

As with other HPC systems, use of the compute nodes on ARCHER is mediated by the PBS job submission system. This is used to ensure that all users get access to their fair share of resources, to make sure that the machine is as efficiently used as possible and to allow users to run jobs without having to be physically logged in.

Whilst it is possible to run interactive jobs (jobs where you log directly into the backend nodes on ARCHER and run your executable there) on ARCHER, and they are useful for debugging and development, they are not ideal for running long and/or large numbers of production jobs as you need to be physically interacting with the system to use them.

The solution to this, and the method that users generally use to run jobs on systems like ARCHER, is to run in *batch* mode. In this case you put the commands you wish to run in a file (called a job script) and the system executes the commands in sequence for you with no need for you to be interacting.

3.10.1 Using PBS job scripts

We will first run the same serial program on the compute nodes. Look at the batch script:

```
input
    emacs -nw sharpen.pbs
```

The first line specifies which *shell* to use to interpret the commands we include in the script. Here we use the Bourne Again SHell (bash), which is the default on most modern systems. The *-login* option tells the shell to behave as if it was an interactive shell.

The line *-l select=[nodes]* is used to request the total number of compute nodes required for your job (1 in the example above).

The #PBS lines provide options to the job submission system where “-l select” specifies that we want to reserve 1 compute node for our job - the minimum job size on ARCHER is 1 node (24 cores); the “-l walltime=00:01:00” sets the maximum job length to 1 minute; “-A y14” sets the budget to charge the job to “y14”; “-N sharpen” sets the job name to “sharpen”.

The remaining lines are the commands to be executed in the job. Here we have a comment beginning with “#”, a directory change to \$PBS_O_WORKDIR (an environment variable that specifies the directory the job was submitted from) and the aprun command (this command tells the system to run the jobs on the compute nodes rather than the frontend nodes).

Jobs can only be run on the compute nodes with the aprun parallel job launcher. The line `aprun -n 1 ./sharpen` runs a single copy of our serial executable.

3.10.2 Submitting scripts to PBS

Simply use the `qsub` command:

```
input
qsub -q RXXXXXX sharpen.pbs
output
58306.sdb
```

You will be given the number of the reserved course queue RXXXXXX. The jobID returned from the `qsub` command is used as part of the names of the output files discussed below and also when you want to delete the job (for example, you have submitted the job by mistake).

3.10.3 Monitoring/deleting your batch job

The PBS command `qstat` can be used to examine the batch queues and see if your job is queued, running or complete. `qstat` on its own will list all the jobs on ARCHER (usually hundreds) so you can use the “-u \$USER” option to only show your jobs:

```
input
qstat -u $USER
output
```

Job ID	Username	Queue	Jobname	SessID	NDS	TSK	Req'd Memory	Req'd Time	Elap S	Time
58306.sdb	user	standard	sharpen	--	1	36	--	00:01	Q	--

“Q” means the job is queued, “R” that it is running and “E” that it has recently completed; if you do not see your job, it usually means that it has completed.

If you want to delete a job, you can use the `qdel` command with the jobID. For example:

```
input
qdel 58306.sdb
```

3.10.4 Finding the output

The PBS system places the output from your job into two files: `<job name>.o<jobID>` and `<job name>.e<jobID>` (note that the files are only produced on completion of the job). The `sharpen.o<jobID>` file contains your job output; `sharpen.e<jobID>` contains any errors.

```
input
cat sharpen.o58306
output
Image sharpening code running in serial
Input file is: fuzzy.pgm
Image size is 564 x 770
Using a filter of size 17 x 17
Reading image file: fuzzy.pgm
... done
Starting calculation ...
On core 0
... finished
Writing output file: sharpened.pgm
... done
Calculation time was 5.400482 seconds
Overall run time was 5.496556 seconds
Application 166570 resources: utime ~6s, stime ~0s, ...
```


4 Running in Parallel

4.1 Running a parallel job on the compute nodes

So far, have run a serial job on a single CPU core and it will probably run slower than if you had compiled and run on your own laptop! Like any supercomputer, the power of ARCHER comes from having many thousands of CPU-cores in the same system.

In this section we will run the sharpen program on many CPU-cores and look at how long it takes to run.

Repeat the same procedure as above but use the parallel MPI version of the code in C-MPI or F-MPI. By default, the batch script is set up to run on 4 cores. You should see that that the parallel code is faster than the serial one.

4.2 Timings

If you examine the log file you will see that it contains two timings: the total time taken by the entire program (including IO) and the time taken solely by the calculation. The image input and output is not parallelised so this is a serial overhead, performed by a single processor. The calculation part is, in theory, perfectly parallel (each processor operates on different parts of the image) so this should get faster on more cores.

You should run a number of jobs and fill in Table 1 below: the IO time is the difference between the calculation time and the overall run time; the total CPU time is the calculation time multiplied by the number of cores.

- If you want to run on more than 24 cores then you will need to request more than one node with the “select=” option to PBS – each ARCHER node only has 24 CPU-cores.

Once you have completed Table 1, you should have the information required to complete Table 2 to compute the speedup of the parallel code overall, of the calculation part and of the IO part. Remember that the speedup is the ratio of runtime at 1 core compared to the runtime at P cores.

Look at your results – do they make sense? Given the structure of the code, you would expect the IO time to be roughly constant, and the performance of the calculation to increase linearly with the number of cores: this would give a roughly constant figure for the total CPU time. Is this what you observe?

4.3 Plotting the speedup

You should also produce a speedup plot where you with number of cores (on the x-axis) versus speedup (on the y-axis). Plot all three speedup curves on the same graph.

To aid you with plotting graphs we provide a small Python plotting utility, plotxy.py that you can download from the ARCHER website with:

```
wget http://www.archer.ac.uk/training/courses/2016/PracIntroHPC/util/plotxy.py
```

If you put your data in CSV (comma-separated value) form in a file then it will plot them for you. For example, to plot the speedup data you would create a file with five values per line: the number of cores followed by each of the speedup values for that core count:

```
<cores>, <ideal>, <overall>, <calculation>, <io>
```

You could create this file using Emacs:

```
emacs -nw speedup.csv
```

and the first couple of lines may look like:

```
1, 1.000, 1.000, 1.000, 1.000
2, 2.000, 1.877, 1.920, 1.014
```

Once you have the data in a CSV format you can plot it with the commands:

```
module load anaconda
python plotxy.py speedup.csv speedup.png
```

(The Anaconda package contains many useful Python libraries, including those that can be used for plotting data: matplotlib). The utility will save a PNG image containing the plot in the second file name you specify on the command line. You can view the plot using the “display” command as you used for the image:

```
display speedup.png
```

This is a very simple plot to allow you to quickly preview results – if you were plotting for including in a report or paper you would produce something more elaborate (including axis labels, proper series labels, etc.).

4.4 OpenMP code

If you are interested, you can try running the OpenMP code in *C-OMP* or *F-OMP*.

Note that to change the number of cores that the program uses you must set the environment variable `OMP_NUM_THREADS`. This is already done for you in the PBS script, although you will need to change the actual number yourself.

Note that you can run using multiple threads on the login nodes – you must set `OMP_NUM_THREADS` before you run, e.g.:

```
export OMP_NUM_THREADS=4
./sharpen
```

The maximum number of cores you can use with OpenMP is 24.

5 Tables

# Cores	Overall run time	Calculation time	IO time	Total CPU time
1				
2				
4				
7				
10				
24				
48				
96				
192				

Table1: Time taken by parallel image processing code

# Cores	Ideal Speedup	Overall speedup	Calculation Speedup	IO Speedup
1				
2				
4				
7				
10				
24				
48				
96				
192				

Table2: Speedup for parallel image processing code