

ARCHER Single Node Optimisation Practical Notes

Getting started

Logging on to ARCHER

You should have been given a guest account ID – referred to generically here as `guestXX` and password.

These credentials can be used to access ARCHER using

```
ssh -X guestXX@login.archer.ac.uk
```

or with the SSH client of your choice (`-X` ensures that graphics are routed back to your desktop). Once you have successfully logged in you will be presented with an interactive command prompt.

For more detailed instructions on connecting to ARCHER, or on how to run commands, please see the Appendix.

Download and extract the exercise files

Firstly, change directory to make sure you are on the `/work` filesystem on ARCHER.

```
guestXX@archer:~> cd /work/y14/y14/guestXX/
```

`/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` filesystem.** ARCHER compute nodes cannot access the `/home` filesystem at all: any jobs attempting to use `/home` will fail with an error.

Use the following commands (on ARCHER) to get the exercise files archive from the web and unpack it:

```
cp /home/z01/shared/SNO.tar .
tar xvf SNO.tar
```

Exercise 1: CrayPat

The code for this exercise is in `ArcherOpt/VH1`. Please follow the separate instructions in `CrayPAT-intro.pdf`, but please note that you need to load the Cray NetCDF module before compiling:

```
module load cray-netcdf
```

Exercise 2: Optimisation

The `ArcherOpt/MD/*` directory, where `*` is either `C` or `Fortran`, contains a sequential implementation of a molecular dynamics simulation which has been deliberately written to have poor performance. Use profiling and compiler listings to look for the performance problems. Once you have exhausted the compiler's ability to optimise, try some code modifications.

Exercise 3: Bandwidth and NUMA

The example code can be found in `ArcherOpt/Stream/*`. This is the well-known STREAM benchmark for measuring memory bandwidth. Use the Makefile to compile the code, and run it using different numbers of threads using the supplied batch script. Does the bandwidth scale linearly with processors? Now try removing the OpenMP loop directive from the initialisation of the arrays. How does the performance change? You can also try using the "wrong" schedule for the loop, or selecting different sets of cores to run on.

Extra exercise

Try reducing the array size `N` by a factor of 100 or 1000 (and increase the repetition count `NTIMES` by the same amount).

Exercise 4: Cache blocking

The `ArcherOpt/Matmul/*` directory contains a simple matrix multiplication code. Try implementing cache blocking by hand, and see the effect on the performance (stick to `-O0` to stop the compiler doing its own optimisation!).

Exercise 5: OpenMP

The `ArcherOpt/MolDyn/*` directory contains a (not very efficient) OpenMP parallel version of a simple molecular dynamics code (not the same as in Exercise 2!). Run the code using the script supplied with the script supplied to measure the performance on 1, 2, 4, 8, 12 and 24 threads.

Try to identify the performance bottlenecks and fix them!

Appendix

Detailed Login Instructions

Procedure for Mac and Linux users

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

```
local$ ssh -X guestXX@login.archer.ac.uk  
Password:
```

you should be prompted to enter your password.

Procedure for Windows users

Windows does not generally have SSH installed by default so some extra work is required. You need to download and install a SSH client application - PuTTY is a good choice:

- <http://www.chiark.greenend.org.uk/~sgtatham/putty/>

When you start PuTTY you should be able to enter the ARCHER login address (login.archer.ac.uk). When you connect you will be prompted for your user ID and password.

Running commands

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

```
guestXX@archer:~> ls  
bin  work
```

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

```
guestXX@archer:~> ls -l  
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*:

```
guestXX@archer:~> man ls  
Man: find all matching manual pages  
* ls (1)  
  ls (1p)  
Man: What manual page do you want?  
Man:
```

In the manual, use the spacebar to move down a page, ‘u’ to move up, and ‘q’ to quit and exit back to the command line.

Using the Emacs text editor

As you do not have access to a windowing environment when using ARCHER, Emacs will be used 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, there are no menus that can be accessed using a mouse).

Start Emacs with the *emacs* command and the name of the file you wish to create. For example:

```
guestXX@archer:~> emacs sharpen_batch.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_batch.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.

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:

```
guestXX@archer:~> cat sharpen_batch.pbs
aprun -n 4 ./sharpen
```

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:

```
guestXX@archer:~> less sharpen.c
```

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.

Hardware

Each node of ARCHER consist of two sockets, each containing a 12-core Intel Ivy Bridge processor.

Compiling

The default compilers are the Cray compilers for Fortran 90 and C. To use the Intel or GNU compilers:

```
guestXX@archer:~> module swap PrgEnv-cray PrgEnv-intel
```

or

```
guestXX@archer:~> module swap PrgEnv-cray PrgEnv-gnu
```

The compiler is always invoked with `ftn` or `cc`, but you will need to modify the flags for the different compilers.

Job Submission

To run codes, you should submit a batch job as follows:

```
qsub -q <resnum> scriptfile.pbs
```

where `resnum` is the reservation number for the session.

You can monitor your jobs status with the `qstat` command, and jobs can be deleted with `qdel`.