

Crib Sheet: ARCHER MPI Exercises

David Henty

1 Logging on

Use your username (`guestXX`) and password to access ARCHER: `ssh -X guestXX@login.archer.ac.uk`

You should then change directory to the `/work/` filesystem: `cd /work/y14/y14/guestXX/`.

You should remove any previous versions of the `MPP-templates` files: `rm -rf MPP-templates*`

2 Obtaining source code

You can obtain `MPP-templates.tar` by using `wget` on ARCHER:

```
wget http://www.archer.ac.uk/training/course-material/
2015/04/MPP_MPI_southampton/Exercises/MPP-templates.tar
```

To obtain this URL, cut and paste it from the Exercise Material section of the course web pages. Now unpack the tar file: `tar xvf MPP-templates.tar` and change directory into it: `cd MPP-templates`

3 Compiling code

You can compile the C and Fortran codes directly:

```
cc -o hello hello.c
ftn -o hello hello.f90
```

or using the supplied Makefiles

```
make -f Makefile_c
make -f Makefile_f90
```

4 Running

First make a copy of the PBS template batch file to match your executable name, e.g. `cp mpibatch.pbs hello.pbs`

Now submit to the batch system: `qsub -q R2769880 hello.pbs`

The batch system will respond with a unique ID for your job: `XXXXXXXX.sdb`.

For Wednesday and Thursday the appropriate queue numbers are R2769890 and R2769892.

You can monitor the progress of your job with `qstat -u guestXX`. The letter "Q" indicates the job is queued, "R" that it is running, and no listing (or "E") that it has finished.

When the job has finished, the output will appear in a file called `hello.pbs.XXXXXXXXX`

To change the number of processes you run on, edit the batch file and change the value of `NPROC`.

5 Interactive Access

The procedure described in Section 4 is the standard way of running all jobs on ARCHER, other than the fact that on this course you are submitting to a specially reserved course queue.

Because the queue is reserved for the guest accounts, your jobs should normally run almost immediately (at least within a couple of minutes). However, if the machine is very busy then there can be a longer delay. In these situations, it may be best for each guest account to reserve a single node from the reservation for several hours; this is called *interactive access* and ensures instant turnaround of jobs.

5.1 Reserving a node

To reserve a node, you submit a special command to the batch system:

```
qsub -q R2769880 -IVl select=1,walltime=1:0:0 -A y14
```

which reserves a node for 1 hour. Note that the third character in the `-IVl` argument is the letter “l” and not the number “1”.

When the node is allocated, you are effectively logged into it and you will get back a command prompt like this:

```
guestXX@mom1:~>
```

Note that **you will now be back in your home directory** so you will need to change to the work filesystem: `cd /work/y14/y14/guestXX/`.

5.2 Running a job

You can now run parallel jobs directly from the command line as you have interactive access to one of the compute nodes. For example, to run on 4 processes:

```
aprun -n 4 ./hello
```

Do not log out from this session! You should leave this interactive window open all the time so you can run jobs whenever you want. You will need to open another window on ARCHER to compile your jobs (see below).

5.3 Compiling

There is a slight complication that you cannot compile from within an interactive session. You should log on to ARCHER as described in Section 1 – remembering **not** to delete the files this time! – and use this window for compilation.