

Notes on CUDA practicals on ARC HTC cluster

Mike Giles

1 The ARC HTC cluster

The layout of the [ARC HTC](#) (high throughput computing) cluster is shown in Figure 1.

The gateway node which sits on the external network is `gateway.arc.ox.ac.uk`. It can be accessed through the SSH command:

```
ssh -X username@gateway.arc.ox.ac.uk
```

but only from machines with registered IP addresses.

The login node which sits on the university network is `htc-login.arc.ox.ac.uk`. This can be accessed from the gateway or from machines on the university network through the SSH command:

```
ssh -X username@htc-login.arc.ox.ac.uk
```

The `-X` option provides X-window forwarding which is convenient when using windows-based editors such as `emacs`.

The nodes `htc-g045` to `htc-g049` each have 8 V100 GPUs. We may also use the `htc-g019` node which has 4 A100 GPUs.

2 ARC User Guide

The ARC User Guide is available at <https://arc-user-guide.readthedocs.io/en/latest/>




Figure 1: ARC HTC cluster with gateway and login nodes

3 CUDA

To use CUDA on ARC HTC you need to enter the command:

```
module load CUDA
```

Amongst other things, this will set your `PATH` and other environment variables so that you can use the NVIDIA compiler `nvcc` and link in the various libraries, and access the relevant header files.

4 Editing and file transfer

`emacs`, `vim` and `nano` are all available on the ARC HTC login node – see next page for info on using `vim` and `nano`.

The alternative is to edit files on your laptop and move them to, and from, ARC using `scp` commands such as:

```
scp local_file username@htc-login.arc.ox.ac.uk:/home/remote_file
scp username@htc-login.arc.ox.ac.uk:/home/remote_file local_file
```

See the next page for SCP applications for Windows users.

5 Doing the practicals

A tar file for all of the practicals can be obtained by using the command:

```
wget https://people.maths.ox.ac.uk/gilesm/cuda/practicals.tar.gz
```

and then untarred using the command:

```
tar zxvf practicals.tar.gz
```

For each practical you compile the code using the supplied Makefile and the command `make`

To run the compiled code as a small batch job on one of the GPU compute nodes you use a command of the form:

```
sbatch job
```

where `job` is the name of a job submission script which is provided for the first few practicals.

You should see the slurm batch management system return something like the following:

```
sbatch: GPU gres requested, checking settings/requirements...
Submitted batch job 3051862
```

Once the job has executed you will see an output from slurm, which in the case of batch job 3051862 will be in the file `slurm-3051862.out`

More information on job submission scripts is given in the [ARC User Guide](#).

6 Links for help on Linux commands and editors

Intro to Linux commands for Windows users

<https://kinsta.com/blog/linux-commands/>

<https://www.geeksforgeeks.org/linux-vs-windows-commands/>

<https://www.hostinger.co.uk/tutorials/linux-commands>

<https://blog.robertelder.org/intro-to-linux-command-line-for-windows-users/>

<https://phoenixnap.com/kb/linux-commands>

nano editor

<https://www.nano-editor.org/>

<https://www.nano-editor.org/dist/latest/cheatsheet.html>

<https://help.ubuntu.com/community/Nano>

vim editor

<https://opensource.com/article/19/3/getting-started-vim>

https://opensource.com/sites/default/files/gated-content/cheat_sheet_vim_final_v2.0.pdf

7 Other tools for Windows users

Some Windows users like [MobaXTerm](#) which combines a remote terminal, file copying and X11 server in one (free) product.

[WinSCP](#) is another popular alternative for copying files to/from remote servers.