Rainer Spurzem

Introduction to GPU Accelerated Computing

Mon-Thu, Feb. 11-14, 2019, Uni Heidelberg, GPU Block Course

1. Getting Access to the kepler computer

Create an ssh key Linux/Unix Users (most things also work inside an open ssh terminal under Windows):

`ssh-keygen -t rsa` (passphrase should NOT be empty)

It produces a private key `id_rsa` and a public key `id_rsa.pub`

Send the public key by email to spurzem@ari.uni-heidelberg.de
After this has been installed on kepler, you can login with

```
ssh lecturenn@kepler.ari.uni-heidelberg.de
```

(lecturenn is your account on the system, such as lecture01, lecture02, …)

If you give a non-standard name to your private key, you need to do this:

```
ssh -i 'full_path_to_private_key' lecturenn@kepler.ari.uni-heidelberg.de
```

If it is e.g. in *home/data/xyz/.ssh/my_id.rsa* you need to use inside the '..' above.

For Windows Users: you can use a terminal window, which works like a Linux command window, compared
https://www.wikihow.com/Use-SSH (check especially 'create encrypted keys and the chmod command, that may cause a problem)
Second Way: Use the putty client program for ssh;
https://putty.org/ (here you need to make sure to copy and paste the openssh public key to send by email!).

2. After login to the kepler computer

```
ls -lt
```

shows you files in your home directory, there should be gpu-lecture.tar.gz and worknb6.tar.gz

```
tar xvfz gpu-lecture.tar.gz
```

creates the subdirectory gpu-lecture:

```
cd gpu-lecture ; ls -ltr
```

shows the course exercises, in several subdirectories (hello, add, ….), next:

```
module load cuda ; module list
```

is needed to make the CUDA and NVIDIA software for GPU available.

3. How to run an exercise

Example of hello world program:

```
cd hello
nvcc -o hello.out hello.c
./hello.out
nvcc -o hello.out hello.cu
./hello.out
nvcc -o device.out device.cu
./device.out
```

produces output, saying there are no GPU devices. This is because the kepler computer has a batch system, we are on the login nodes, and not on the 12 worker nodes. Only the worker nodes have GPUs.

We need to use the slurm batch system to submit jobs to the worker nodes; the submission script is provided in gpu_script.sh :

```
sbatch < gpu_script.sh
```

Information about the job and queue status (see also kepler Manual link on our course webpage):

```
squeue
sinfo
```

The result of the batch job will appear on job.nnnn.out and job.nnnn.err ; look for it with

```
ls –ltr
```

For reading and editing files you can use the Linux programs:

```
cat  ; more ; vi ; …?
```

Transfer data files to your own computer:

```
scp lecturenn@kepler.ari.uni-heidelberg.de:gpu-
lecture/hello/job.2345.out .
```

(Note the .)