Rainer Spurzem, spurzem@ari.uni-heidelberg.de

Introduction to GPU Accelerated Computing

Mon-Fri, Mar 1-5, 2021, Uni Heidelberg, GPU Block Course

Tutors:

Katja Stock katja.stock@uni-heidelberg.de
Peter Lysakowski peter.lysakovski@uni-heidelberg.de
Albrecht Kamlah albrechtk@hotmail.de

1. zoom and chat access, general informations

Lecture: Rainer Spurzem
https://zoom.us/j/92682356846?pwd=L1JOZ2s0ekg2SW5DY2c5bmxjb2xYQT09
Meeting ID: 926 8235 6846 ; Passcode: 551062

Group 1: Zoom Tutorium-1, Katja Stock
https://zoom.us/j/91799498402
Meeting ID: 917 9949 8402 ; Passcode: 768181

Group 2: Zoom Tutorium-2 ; Peter Lysakovski
https://zoom.us/j/98446846876
Meeting ID: 984 4684 6876 ; Passcode: 731394

Group 3: Zoom Seminar-2 ; Albrecht Kamlah
Mar 1-3:
https://zoom.us/j/93694063571?pwd=YzdZWmtJRGMyemd5bk5SQy9uakJnZz09
Meeting ID: 936 9406 3571 ; Passcode: sf.hd.2020
Mar 4:
https://zoom.us/j/5699021721?pwd=NUppdjVJVC9UQ09jVjlFSDI5dmdQQT09
Meeting ID: 569 902 1721 ; Passcode: sf.hd.2020

Time in lecture: Mon-Fri 10:15 – 13:00 ; 14:15 – 15:00 ***
Time in Tutorials: Mon-Thu 15:00 – 17:00 ***

****exact times of transition from lecture to tutorials may vary!*

There is a rocket chat open:
https://uebungen.physik.uni-heidelberg.de/chat/group/ws20-gpu-cuda-course

Also for each group 1,2,3:
https://uebungen.physik.uni-heidelberg.de/chat/group/ws20-gpu-cuda-course-1 , -
https://uebungen.physik.uni-heidelberg.de/chat/group/ws20-gpu-cuda-course-2
https://uebungen.physik.uni-heidelberg.de/chat/group/ws20-gpu-cuda-course-3

Rocket Chat: Messages independent of zoom!
Messages are permanent!
In zoom: raise hand (preferred) and chat, but not permanent!

Need to know: This course is not graded!
Certificate of successful participation only! (2 ECTS points)

Necessary condition for certificate:
Regular participation in course, all Mon-Wed exercises done on the computer! The last exercise (Thu) is a homework.
You can turn it in later, talk to your tutor!

This is NOT a programming course!
You should have some knowledge of a higher level programming language (like python, c, c++, fortran, …). We will learn special CUDA extensions of the c-language for programming GPUs (graphical processing unit). This is called GPGPU – general purpose GPU programming.

This is NOT a programming course!
You learn CUDA, but our learning here is by using examples, running them, explaining the background in the lecture, and only make small changes to the program, if any.

Data Handling!
For the homework you produce some simple data, and the task is to plot them. You can use anything you like (jupyter notebook, gnuplot, idl, …); we try to help – but again – this is not part of the course.

2. Getting access to the kepler computer

Create an ssh key Linux/Unix Users (also possible for Windows/Mac):

`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@kepler2.zah.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!).

3. After login to the kepler computer

```
ls -lt
```

shows you files in your home directory, there should be
`gpu-course-master.tgz` and **`gworknb6.tar.gz`** (later)

```
tar xvfz gpu-course-master.tgz
```

creates the subdirectory gpu-course:

```
cd gpu-course ; ls -ltr
```

shows the course exercises, in several subdirectories (0_hello, 1_add, …), next:

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

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

4. 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 ; nano ; pico ; …?
```

Transfer data files to your own computer:

```
scp  lecturenn@kepler.ari.uni-heidelberg.de:gpu-course/0_hello/job.2345.out .
```

(Note the .)