

# Getting started with CUDA

## Forewords

---

Edwin Carlinet, Joseph Chazalon {firstname.lastname@lrde.epita.fr}

April 2020

EPITA Research & Development Laboratory (LRDE)



Much of this lesson is based on these great resources — *Look them up to go further!*

- The course “GPU Teaching Kit”, 2019, licensed by NVidia and the University of Illinois under the Creative Commons Attribution-Non Commercial 4.0 International License.
- The book “Programming massively parallel processors” (Third Edition), D. Kirk and W. Hwu, Elsevier, 2017.
- The manual “CUDA C Programming Guide”, NVidia, v10.1.243 (August 19, 2019).



*This lesson is licensed by E. Carlinet and J. Chazalon under the Creative Commons Attribution-Non Commercial 4.0 International License.*

## Important remarks about this lesson

The format of this second lecture is **different from** the ones of the first and third lectures:

1. We first give you all the training material: slides, videos and instructions for the practice session
2. The material is divided into the following sections, to ease browsing:
  - CUDA overview
  - Host view of GPU computation
  - Compilation and Runtime
  - Kernel programming
3. For each section, you have work by yourself to:
  - Read/View the available material
  - Work on the related practice session
4. The time slot for the 2nd lecture (Tue. afternoon) is maintained but your will work alone
5. The time slot for the 1st practice session (Wed. afternoon) will be a live chat to answer all your questions: be ready and work before!

## Using EPITA's CUDA GPUs remotely

Under the current exceptional circumstances, you are provided with a remote SSH access to the machines of room SM14 which are equipped with NVidia GTX 1050 cards:

- Each student has access to a remote machine.
- A machine is shared by two students.

With the listing we provided you with, use your login and the port number we assigned to you to connect to a given machine:

```
ssh -p $PORT $LOGIN@cri.epita.fr
```

You will need your **CRI password**.

*Please note that those machines do not have an X server, so you will have to use the terminal only.*

## Using your own computer

You can use your own computer but we will provide **no assistance** for CUDA installation.

Indeed, this is a tedious task on some systems / distributions.

If you choose this path, make sure you are ready to:

- mess your OS installation;
- spend two days doing this;
- eventually fail.

If you manage to compile and launch a basic CUDA *hello world*, do not touch anything no matter the CUDA version you have: this should be sufficient for most of the work of this course.