Getting started with CUDA
Forewords

Edwin Carlinet, Joseph Chazalon {firstname.lastname@epita.fr}
Fall 2023
EPITA Research Laboratory (LRE)


Slides generated on September 8, 2023

Sources and license

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).

Using CUDA-ready hardware

2 options

You should either:

- use **OpenStack machines** (recommended option, GeForce GTX 1650)
- use your personal computer (no assistance provided)

OpenStack machines

We have **5 VMs with GPU cards** which are **available 24/7** (running NixOS).

You can use them during practice sessions, and you should use them for your projects.

You may need to configure your SSH keys on https://cri.epita.fr/.
PRO TIP: copy-paste the content of `id_*.pub`, not `id_*`

To connect, use:

```
ssh -X -p $PORT ${EPITA_LOGIN}@gpgpu.image.lrde.iaas.epita.fr
```

- Replace `$PORT` with a random value between 22000 and 22004.
- Replace `${EPITA_LOGIN}` by your login.
- You may need to type your password once to get a Kerberos ticket (not sure about this).

To access your AFS storage, use `kinit $LOGIN` and `aklog`.
Make sure you replace the first dot of your login (for recent ones) by and underscore (_).

Example for `alex.login` :

```
kinit alex_login
aklog
```


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 up your OS installation;
- spend two days trying to get this work;
- eventually fail.

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

Using NVidia tools on NixOS

Please refer to **our online documentation** on the course page.
https://www.lrde.epita.fr/~carlinet/cours/GPGPU/#how-to-use-openstack-gpus

In particular, we provide you with a Nix Shell configuration file which enables all the tools, in a virtual environment-like fashion:

```
wget https://www.lrde.epita.fr/~carlinet/cours/GPGPU/shell.nix
nix-shell shell.nix
```

Note that you will have to reactivate this environment for **every new shell you launch**.