HPC.NRW

PPCES 2022: GPU / ML / DL Exercise Instructions

HPC.NRW Competence Network

Course Organization

HPC.NRW Competence Network

PPCES 2022: How to start the exercise environment

Background: Connection to RWTH Cluster

HPC.NRW

- Connect to login18-x-1 or login18-x-2
- For this exercise we need
  - Command line (Start a job on our cluster)
  - Firefox Browser (Jupyter Lab)

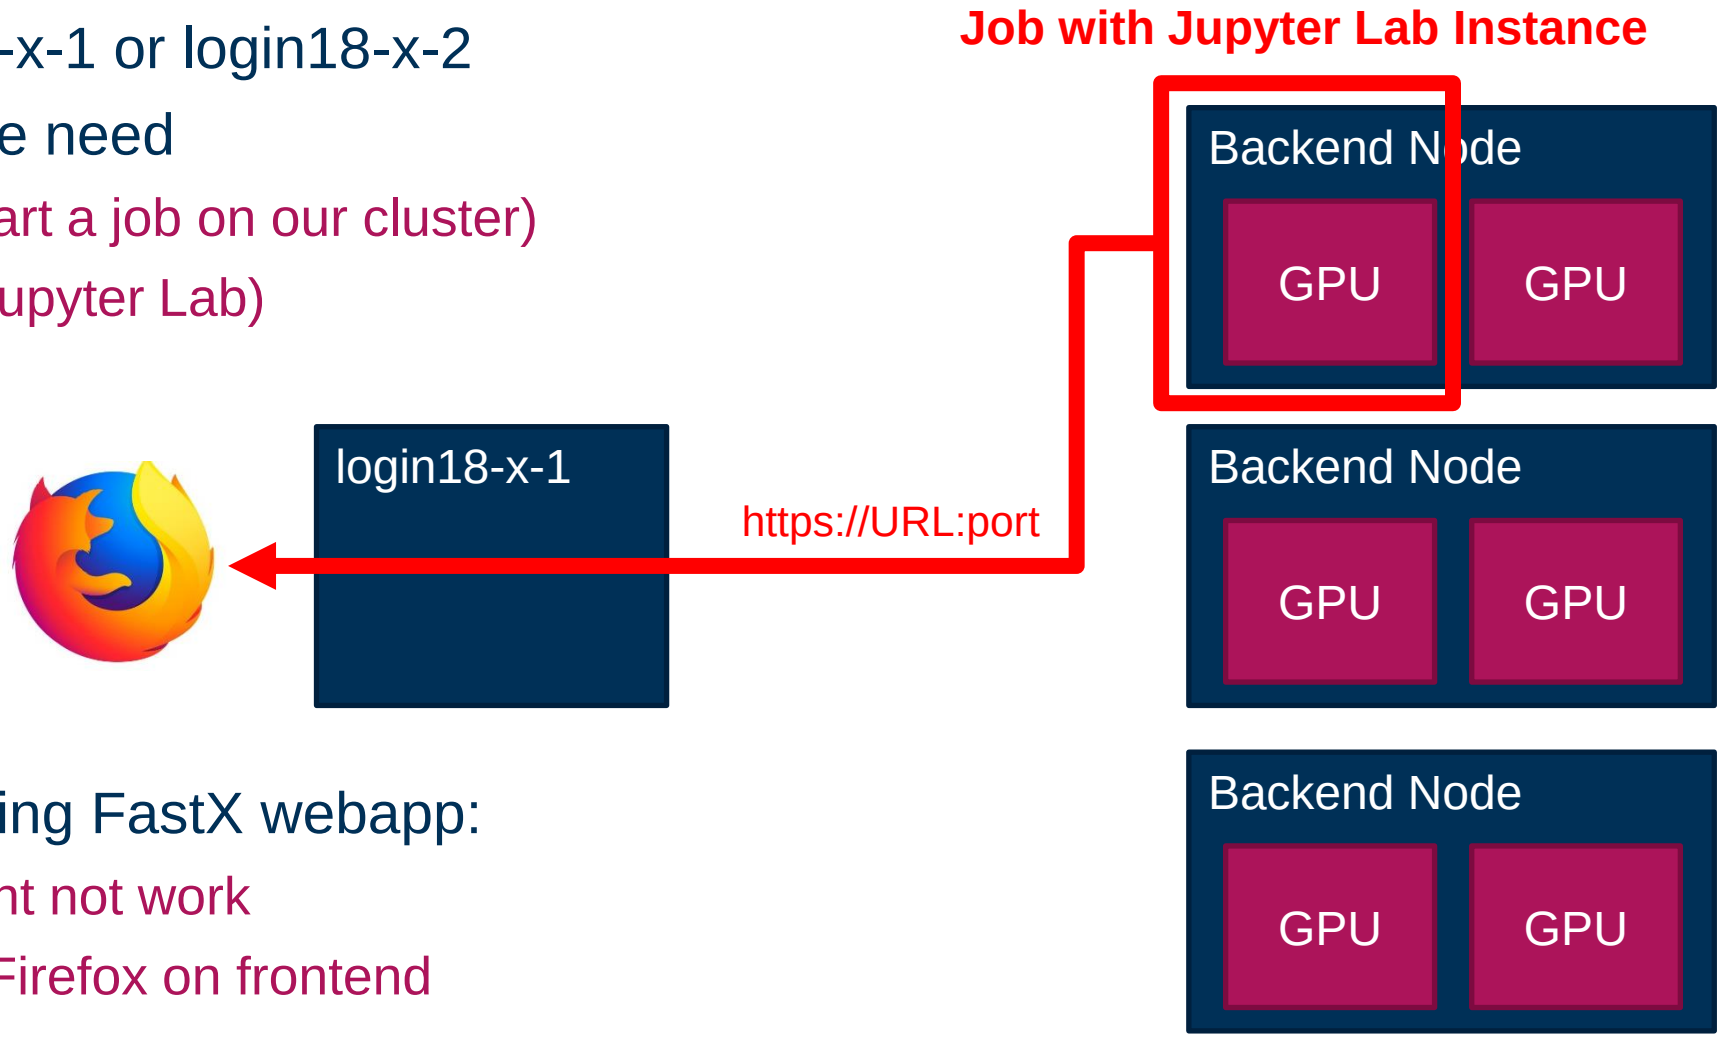

- **Note:** If you are using FastX webapp:
  - Copy & Paste might not work
  - Open PDF inside Firefox on frontend

Procedure / Steps

HPC.NRW

1. Download the course material and exercises

2. Request and start a new interactive job
3. Run a Singularity container & Jupyter Lab
4. Connect to Jupyter Lab
5. Close the browser & stop Jupyter Lab, container and job

Singularity containers containing:

scikit-learn

TensorFlow + Horovod

6. Logout

1. Download the course material and exercises

HPC.NRW

– Start a terminal session

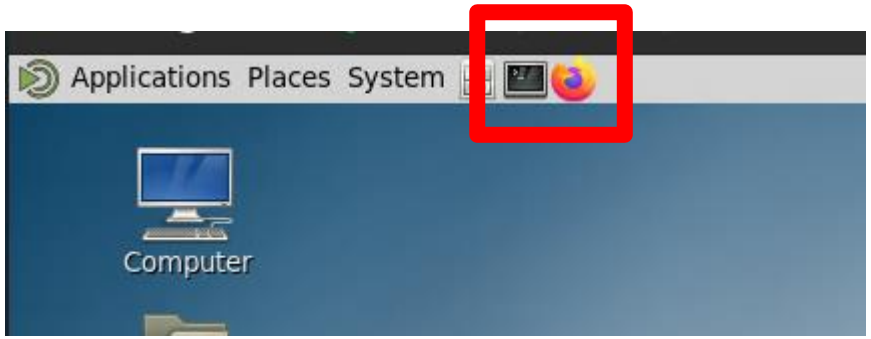

– Download the course material

```
# Change to the directory where you want to save your material
cd <working directory>

# Download this slide deck
wget https://blog.rwth-aachen.de/itc-events/files/2022/02/2022-ppces-ML-DL-instructions.pdf

# Download exercises & scripts
wget https://blog.rwth-aachen.de/itc-events/files/2022/02/2022-ppces-excercises-ml-dl.tar.gz

# Unpack the exercises
tar -xzvf 2022-ppces-excercises-ml-dl.tar.gz
```

Predefined Scripts for Steps 2 & 3

HPC.NRW

2. Request and start a new interactive job using SLURM

3. Run a Singularity container & Jupyter Lab

– For **scikit-learn** exercises:

```
# start interactive job and container
zsh batch-scikit.sh
```

– For **TensorFlow** exercises:

```
# start interactive job and container
zsh batch-tensorflow.sh
```

– For **Horovod** exercises:

```
# start interactive job and container
zsh batch-horovod.sh
```

Predefined Scripts for Steps 2 & 3 (Example: Tensorflow)

HPC.NRW

- What is happening inside?

```
srun    --time=01:00:00                 \
        --partition=c18g                \
        --gres=gpu:volta:1              \
        --reservation=ppces_gpu_22      \
        ...                             \
        zsh ./jupyter-tensorflow.sh
```

- **Note:** You might wait a bit to get a node assigned depending on the load on the system

```
srun: [I] No output file given, set to: output_%j.txt
srun: job 20040570 queued and waiting for resources
```

- Content of `jupyter-tensorflow.sh`
  - Load required modules for NVIDIA GPUs (cuda and cudnn)
  - Start the desired Singularity container + map the exercise directory
  - Within that container start Jupyter Lab

4. Connect to Jupyter Lab

HPC.NRW

– After launching Jupyter Lab you will see something like

```
[I 2021-03-23 17:32:08.087 ServerApp] Use Control-C to stop this server and shut down all kernels (twice to skip confirmation).
[W 2021-03-23 17:32:08.179 ServerApp] No web browser found: could not locate runnable browser.
[C 2021-03-23 17:32:08.180 ServerApp]

    To access the server, open this file in a browser:
        file:///project/.local/share/jupyter/runtime/jpserver-231459-open.html
    Or copy and paste one of these URLs:
        http://ncm0161.hpc.itc.rwth-aachen.de:8888/lab?token=1e6362b6a777c8420daf678462c9f36092518eabd5af329d
     or http://127.0.0.1:8888/lab?token=1e6362b6a777c8420daf678462c9f36092518eabd5af329d
```

– Open the link (here: http://ncm0161...) with Firefox

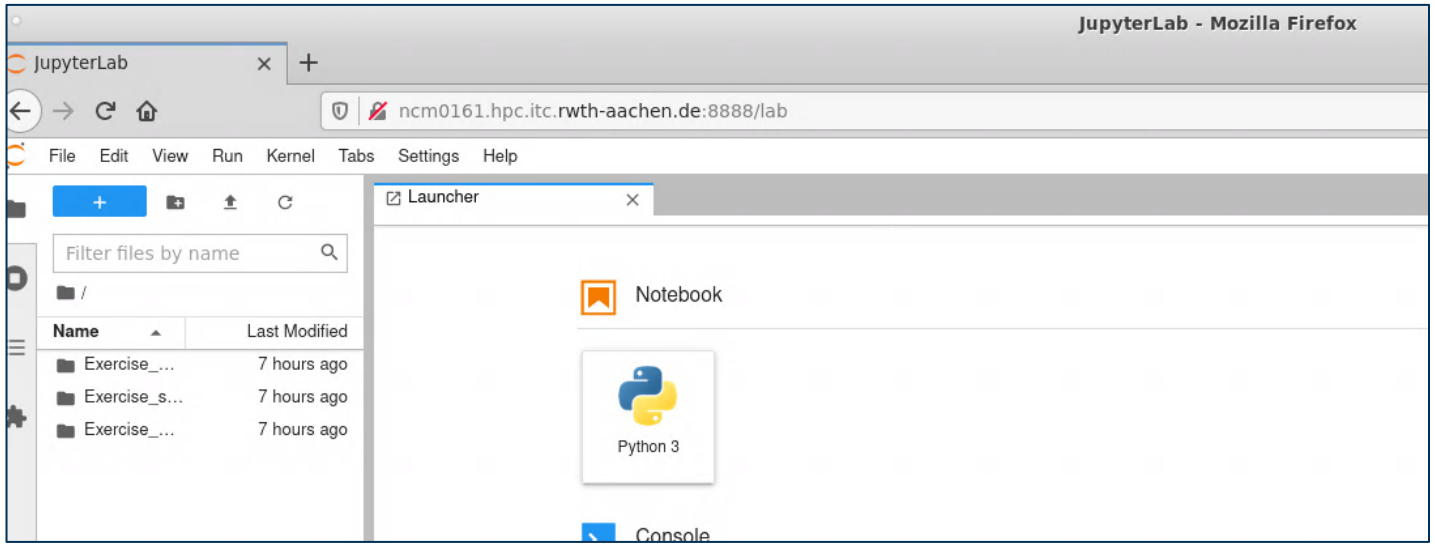

Note:
Jupyter Lab might output several infos / warnings on the command line.

Ignore it as far as everything is working in the browser

5. Close the browser & stop Jupyter Lab and container

HPC.NRW

– Open the terminal where you started Jupyter Lab

```
[I 2021-03-23 17:32:08.087 ServerApp] Use Control-C to stop this server and shut down all kernels (twice to skip confirmation).
[W 2021-03-23 17:32:08.179 ServerApp] No web browser found: could not locate runnable browser.
[C 2021-03-23 17:32:08.180 ServerApp]

    To access the server, open this file in a browser:
        file:///project/.local/share/jupyter/runtime/jpserver-231459-open.html
    Or copy and paste one of these URLs:
        http://ncm0161.hpc.itc.rwth-aachen.de:8888/lab?token=1e6362b6a777c8420daf678462c9f36092518eabd5af329d
     or http://127.0.0.1:8888/lab?token=1e6362b6a777c8420daf678462c9f36092518eabd5af329d
```

– Press **CTRL + C** to stop Jupyter Lab, container and the job
  – **Note:** You might need to confirm with “Y” or press CTRL + C several times

6. Logout

- If you are done with the exercises
  - Log out of the cluster front end
  - Otherwise, session will remain active

Course Organization

HPC.NRW Competence Network

PPCES 2022: Troubleshooting

Troubleshooting

HPC.NRW

- **Problem:** Images not shown in Jupyter Lab

```
plt.axis('off')
plt.title('Quantized image (64 colors, Random)')
plt.imshow(recreate_image(codebook_random, labels_random, w, h))
plt.show()
```

```
Automatically created module for IPython interactive environment
Fitting model on a small sub-sample of the data
done in 0.267s.
Predicting color indices on the full image (k-means)
done in 0.180s.
Predicting color indices on the full image (random)
done in 0.790s.
<Figure size 640x480 with 1 Axes>
<Figure size 640x480 with 1 Axes>
<Figure size 640x480 with 1 Axes>
```

- **Solution:**
  - Select cell again
  - Execute it again using the “Run” button

```
Predicting color indices on the full image (random)
done in 0.731s.
```

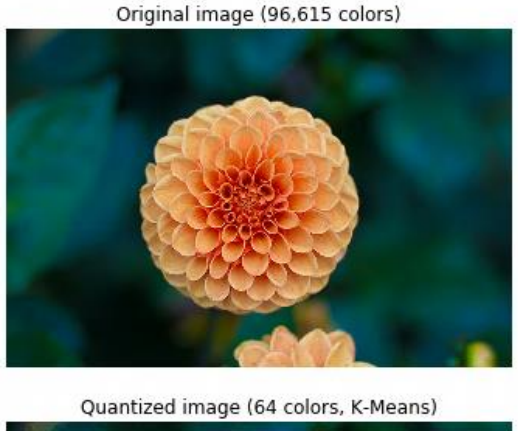

Troubleshooting

HPC.NRW

- **Problem:** “InternalError: CUDA runtime implicit initialization on GPU:0 failed. Status: all CUDA-capable devices are busy or unavailable”

- **Reasons:**
  - You are either working on a GPU frontend
    - All available GPUs are currently in use
    - GPUs on frontends are configured in “`Exclusive_Process`” computation mode
  - Jupyter kernels might still be running on a GPU
    - Solution: stop your running kernels

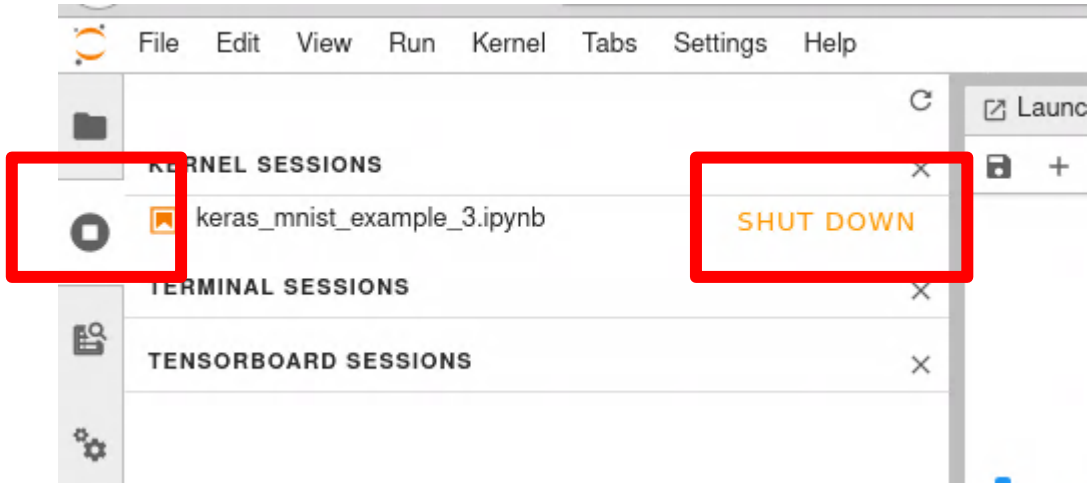

Troubleshooting

HPC.NRW

– **Problem:** Tensorboard is not working

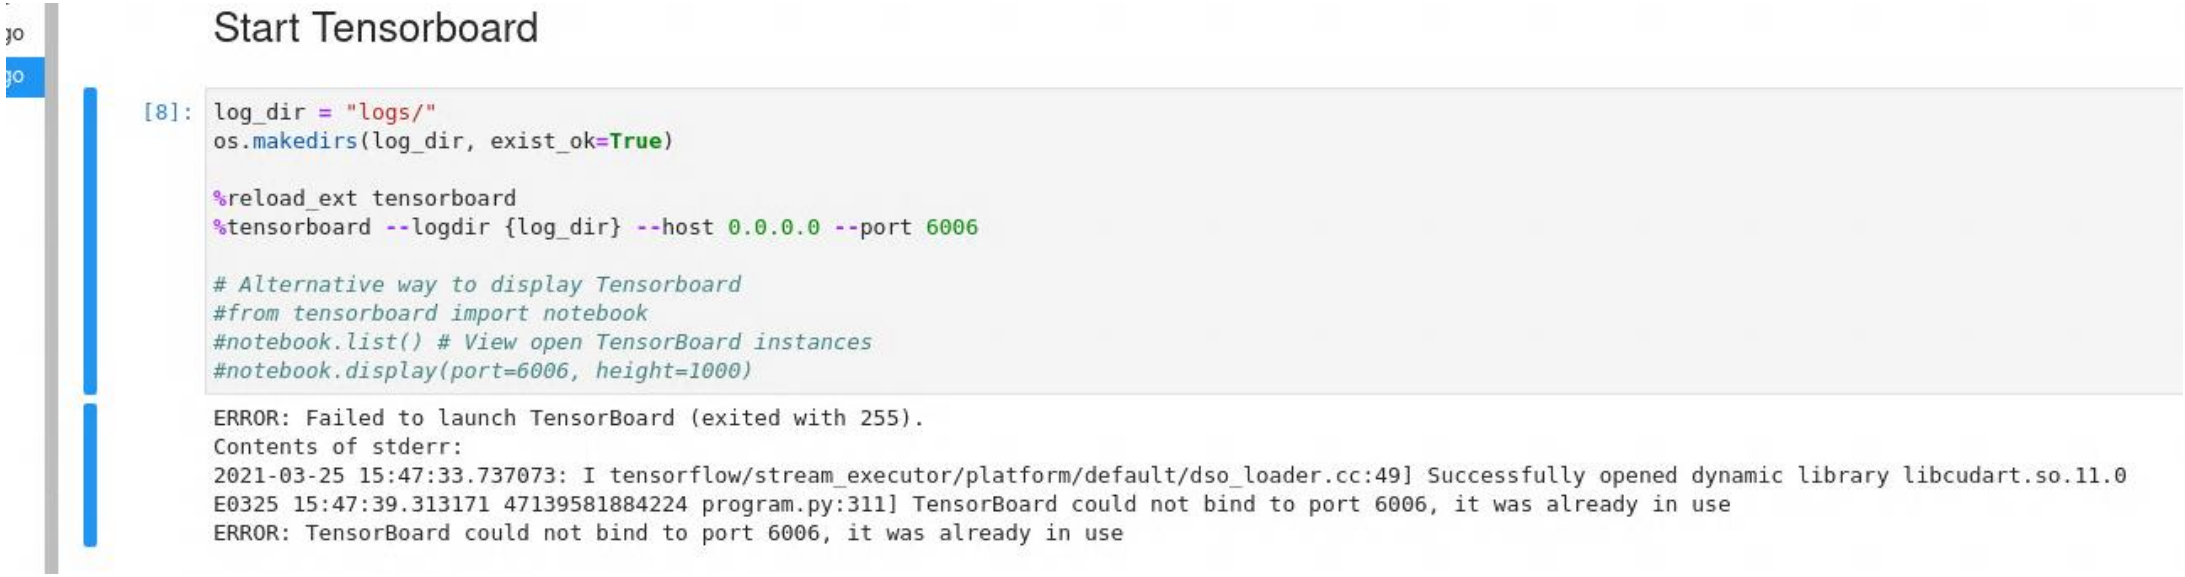

– **Reason:** There is already some instance or program using that port
  – Don’t worry. This will not stop you from testing the rest of the exercises
  – Just change the port in the script to a different one