Bioinformatics Lab: Introduction of High Performance Computing

Ivan Gesteira Costa & Zhijian Li
Institute for Computational Genomics

Institute for Computational Genomics
RWTH AACHEN UNIVERSITY

How to use

You have to use the *secure shell protocol* (**ssh**) to log in

$ ssh <username>@login-g.hpc.itc.rwth-aachen.de

How to use

See what is running: *nvidia-smi*

```
Sat May 26 14:30:06 2018
+-----------------------------------------------------------------------------+
| NVIDIA-SMI 384.81                 Driver Version: 384.81                    |
|-------------------------------+----------------------+----------------------+
| GPU  Name        Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|         Memory-Usage | GPU-Util  Compute M. |
|===============================+======================+======================|
|   0  Tesla P100-SXM2...  Off  | 00000000:06:00.0 Off |                    0 |
| N/A   45C    P0    42W / 300W |      0MiB / 16276MiB |      0%   E. Process |
+-------------------------------+----------------------+----------------------+
|   1  Tesla P100-SXM2...  Off  | 00000000:84:00.0 Off |                    0 |
| N/A   46C    P0    37W / 300W |      0MiB / 16276MiB |      2%   E. Process |
+-------------------------------+----------------------+----------------------+

+-----------------------------------------------------------------------------+
| Processes:                                                       GPU Memory |
|  GPU       PID   Type   Process name                             Usage      |
|=============================================================================|
|  No running processes found                                                 |
+-----------------------------------------------------------------------------+
```

How to use

See what is running: *nvidia-smi*

GPU ID + type

```
Sa                                    6 2018
+-----------------------------------------------------------------------------+
|                                1                 Driver Version: 384.81     |
|-------------------------------+----------------------+----------------------+
| GPU  Name        Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|         Memory-Usage | GPU-Util  Compute M. |
|===============================+======================+======================|
|   0  Tesla P100-SXM2...  Off  | 00000000:06:00.0 Off |                    0 |
| N/A   45C    P0    42W / 300W |      0MiB / 16276MiB |      0%   E. Process |
+-------------------------------+----------------------+----------------------+
|   1  Tesla P100-SXM2...  Off  | 00000000:84:00.0 Off |                    0 |
| N/A   46C    P0    37W / 300W |      0MiB / 16276MiB |      2%   E. Process |
+-------------------------------+----------------------+----------------------+

+-----------------------------------------------------------------------------+
| Processes:                                                       GPU Memory |
|  GPU       PID   Type   Process name                             Usage      |
|=============================================================================|
|  No running processes found                                                 |
+-----------------------------------------------------------------------------+
```

How to use

See what is running: *nvidia-smi*

GPU ID + type

GPU memory

```
Sa                                    6 2018
+-----------------------------------------------------------------------------+
|                                1                 Driver Version: 384.81     |
|-------------------------------+----------------------+----------------------+
| GPU  Name        Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|         Memory-Usage | GPU-Util  Compute M. |
|===============================+======================+======================|
|   0  Tesla P100-SXM2...  Off  | 00000000:06:00.0 Off |                    0 |
| N/A   45C    P0    42W / 300W |      0MiB / 16276MiB |      0%   E. Process |
+-------------------------------+----------------------+----------------------+
|   1  Tesla P100-SXM2...  Off  | 00000000:84:00.0 Off |                    0 |
| N/A   46C    P0    37W / 300W |      0MiB / 16276MiB |      2%   E. Process |
+-------------------------------+----------------------+----------------------+

+-----------------------------------------------------------------------------+
| Processes:                                                       GPU Memory |
|  GPU       PID   Type   Process name                             Usage      |
|=============================================================================|
|  No running processes found                                                 |
+-----------------------------------------------------------------------------+
```

How to use

See what is running: *nvidia-smi*

GPU ID + type

GPU memory

```
Sa                                    6 2018
+-----------------------------------------------------------------------------+
|                                   1                 Driver Version: 384.81    |
|-------------------------------+----------------------+----------------------+
| GPU  Name        Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|         Memory-Usage | GPU-Util  Compute M. |
|===============================+======================+======================|
|   0  Tesla P100-SXM2...  Off  | 00000000:06:00.0 Off |                    0 |
| N/A   45C    P0    42W / 300W |      0MiB / 16276MiB |      0%   E. Process |
+-------------------------------+----------------------+----------------------+
|   1  Tesla P100-SXM2...  Off  | 00000000:84:00.0 Off |                    0 |
| N/A   46C    P0    37W / 300W |      0MiB / 16276MiB |      2%   E. Process |
+-------------------------------+----------------------+----------------------+

+-----------------------------------------------------------------------------+
| Processes:                                                       GPU Memory |
|  GPU       PID   Type   Process name                             Usage      |
|=============================================================================|
|  No running processes found                                                 |
+-----------------------------------------------------------------------------+
```

Compute model: 1 person

Train CNN on GPU

Install Tensorflow and Keras

```
pip install --user tensorflow-gpu
Pip install --user keras
```

Train CNN on GPU

Set environment variable

```
export PATH=/usr/local_rwth/sw/cuda/9.0.176/bin:$PATH
export LD_LIBRARY_PATH=/usr/local_rwth/sw/cuda/9.0.176/
lib64:$LD_LIBRARY_PATH
```

Load modules

```
module load cuda/90
module load cudnn/7.0.5
```

Download example

```
git clone git@github.com:keras-team/keras.git
```

Train CNN on GPU

Run the example

```
python keras/examples/mnist_cnn.py
```

You can also select the GPU

```
CUDA_VISIBLE_DEVICES=1 python keras/examples/mnist_cnn.py
```

Using the batch system

For short test runs (max. 2 hours) in batch mode:

```
#BSUB -a gpu
```

For long runs (only nighttime and weekend):

```
#BSUB -q gpu
```

Submit your jobs:

```
bsub < MyGPUScript.zsh
```

Using the batch system

#nodes	node names	GPU architecture	GPU type	# GPU-cards	operating time "-a gpu"	operating time "-q gpu"
1	linuxnvc02	Kepler	Tesla K20Xm	2	whole day (short test runs only)	working days: 8:15 pm - 7:30 am weekends: whole day
1	linuxnvc03	Kepler	Tesla K40c	1	whole day (short test runs only)	working days: 8:15 pm - 7:30 am weekends: whole day
9	lng01..lng09	Pascal	P100 SXM2	2	whole day	
2	lns07..lns08	Pascal	P100 SXM2	1	whole day	

Using the batch system

#nodes	node names	GPU architecture	GPU type	# GPU-cards	operating time "-a gpu"	operating time "-q gpu"
1	linuxnvc02	Kepler	Tesla	2	whole day	working days: 8:15 pm - 7:30 am weekends: whole day
1	linuxnvc03	Kepler	Tesla K40c	2	whole day (short test runs only)	working days: 8:15 pm - 7:30 am weekends: whole day
9	lng01..lng09	Pascal	P100 SXM2	2	whole day	
2	lns07..lns08	Pascal	P100 SXM2	1	whole day	

Need a project to use these GPUs

More information

https://doc.itc.rwth-aachen.de/display/CC/GPU+cluster