Practical Example: High Performance Computing

Ivan G. Costa & Martin Manolov

Institute for Computational Genomics

Joint Research Centre for Computational Biomedicine

RWTH Aachen University, Germany

Institute for Computational Genomics
01011011010
10100100101

RWTH AACHEN UNIVERSITY

The Operating System UNIX

- UNIX Development
  - 60s: MULTICS: MULTiplexed Informaion and Computing Service
  - Developed by MIT, Bell Labs and General Electrics
  - Written in Programming Language 1 (PL1) and Assembly

  - Thompson and Ritchie later create UNIX, written in C

Linux – Basic Commands (Revisit)

- command parameter1 parameter2 …

- *man* – **man**ual. Displays manuals for linux packages *(*try *man man). Displays useful information about how to use the package.*

- *pwd* – **p**rint **w**orking **d**irectory. Displays the current working directory

Linux – Basic Commands (Revisit)

- *ls* – list. Displays the content of a directory
  - ls -l ~/Pictures

```
-rw-r--r--   1  martin  martin  101685  Apr 29  13:23  Selection_001.png
-rw-r--r--   1  martin  martin   39835  Apr 29  14:37  Selection_002.png
-rw-r--r--   1  martin  martin   76985  Apr 29  14:48  Selection_003.png
-rw-r--r--   1  martin  martin   27096  May  2  14:04  Selection_004.png
-rw-r--r--   1  martin  martin   62783  May  3  11:58  Selection_005.png
-rw-r--r--   1  martin  martin   65201  May  6  16:11  Selection_006.png
drwxrwxr-x   2  martin  martin    4096  Apr 24  11:41  Temp
```

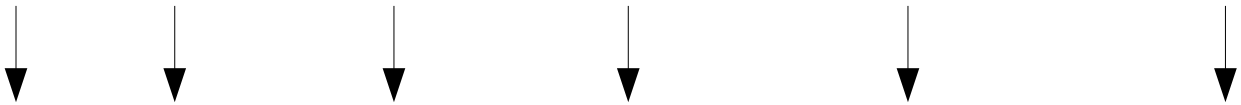

Access rights	Links	Owners/ Groups	Size	Last Modified	Filename

Linux – Basic Commands (Revisit)

- *ls* – **ls**t. Displays the content of a directory

```
-rw-r--r--    1    martin    martin    101685    Apr 29    13:23    Selection_001.png
```

  - Access rights:
    - Type: (**d**)irectory, (**l**)ink, (**-**) a file
    - Rights: (**r**)ead, (**w**)rite, e(**x**)ecute
    - Sequence: *user group anybody*
  - Owner/Group:
    - Owner – first column: *martin*
    - Group – second column: *martin*

- Either can be changed with *chmod* and *chgrp* respectively

- Run executable with full path or from directory with "*./file*"

Linux – Basic Commands (Revisit)

- *cd* – **c**hange **d**irectory. Switches to a new directory, supplied as a parameter

- *mkdir* – **m**a**k**e **dir**ectory. Creates a new directory

- *rm* – **r**e**m**ove. Removes a specified file or directory (“-r”)

- *cp* – **c**o**p**y. Copy a file or directory (“-r”)
- *scp* – **s**ecure **c**o**p**y. Copy a file to or from a remote source
- *rsync* – a fast copying tool. Also works for remote copy

Linux – Basic Commands (Revisit)

- *ln* – **l**i**n**k. Link a file into a new directory
- *echo* – Prints a string to standard output
- *cat* – con**cat**enate. The content of a file is printed to standard output
- *wc* – **w**ord **c**ount. Counts the number of words, rows (“-l”) pr characters (“-c”) in a file.

Linux – Basic Commands (Revisit)

- pipe (“|”) - connect commands

- output (“>”) and input (“<”) for a specific program. Output can also concatinate to existing content, without deleting (“>>”)

Linux – Environment Variables

Variable	Description
PATH	Colon separated list of directories, which will be searched through when entering a name of executable
HOME	The pathname of the home directory.
SHELL	The currently used shell program
USER	The current username

```
$ echo $PATH
/usr/local/sbin

$ export PATH=~/.local/bin:$PATH
$ echo $PATH
/home/martin/.local/bin:/usr/local/sbin
```

Linux – Vim

- A further development of Vi
- Open-Source editor for use inside the Terminal
- Open file with “$ vim filename”
  - Type “i” for insert (edit)
  - ESC to go back to entry menu
  - :q – exit without saving
  - :wq – save and exit
  - :q! - force exit without saving

Linux – Bash Scripts – Exercise

```
#!/bin/bash
echo Download some sample data
wget https://www.costalab.org/wp-content/uploads/2019/05/data2.zip

echo Unzipping the data
unzip data2.zip

echo Removing the archive
rm data2.zip

echo Data ready to work with
```

- Write in a file with .sh – extension
- Change the file mods bits
  - chmod 770 somefile.sh
- Execute the file with "*./somefile.sh*"

RWTH Compute Cluster

- High Performance Computing Infrastructure
- Linux-Based
- Highly Parallelizable
- All RWTH - affiliates granted access, but with limited resource allocation
- More information on: https://doc.itc.rwth-aachen.de/display/CC/General

Source: http://www.itc.rwth-aachen.de

RWTH Selfservice

- Use Selfservice (https://idm.rwth-aachen.de/selfservice/) to manage university accounts and create HPC-account to user the Cluster

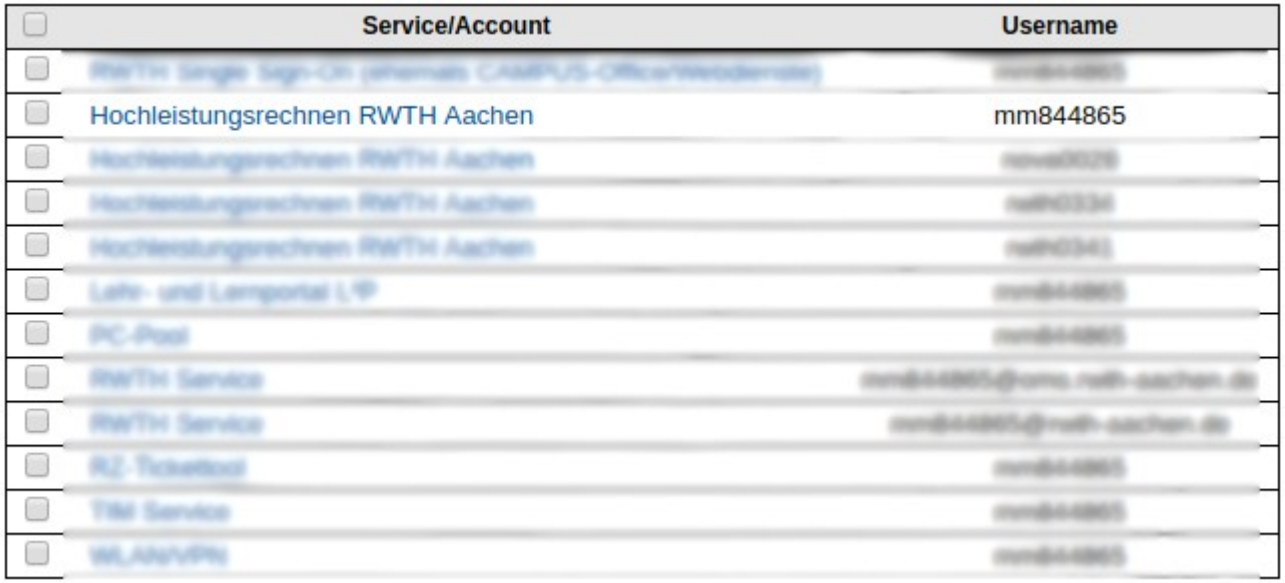

Service/Account	Username
[illegible]	[illegible]
Hochleistungsrechnen RWTH Aachen	mm844865
[illegible]	[illegible]
[illegible]	[illegible]
[illegible]	[illegible]
[illegible]	[illegible]
[illegible]	[illegible]
[illegible]	[illegible]
[illegible]	[illegible]
[illegible]	[illegible]
[illegible]	[illegible]
[illegible]	[illegible]

SSH and Cluster login

- Use established *secure shell protocol (ssh)* connection to log in to the front-end cluster nodes.

  $ ssh <Username>@<Servername(or IP)>

SSH and Cluster login

Server-name	OS	Purpose
login18-1.hpc.itc.rwth-aachen.de	Linux (CentOS 7)	Front-End Dialogue System for CLAIX2018
login18-2.hpc.itc.rwth-aachen.de	Linux (CentOS 7)	Front-End Dialogue System for CLAIX2018
login18-3.hpc.itc.rwth-aachen.de	Linux (CentOS 7)	Front-End Dialogue System for CLAIX2018
login18-4.hpc.itc.rwth-aachen.de	Linux (CentOS 7)	Front-End Dialogue System for CLAIX2018
login.hpc.itc.rwth-aachen.de	Linux (CentOS 7)	Front-End Dialogue System for CLAIX2016
login2.hpc.itc.rwth-aachen.de	Linux (CentOS 7)	Front-End Dialogue System for CLAIX2016
login18-g-1.hpc.itc.rwth-aachen.de	Linux (CentOS 7)	GPU-System (2018)
login18-g-2.hpc.itc.rwth-aachen.de	Linux (CentOS 7)	GPU-System (2018)
login-g.hpc.itc.rwth-aachen.de	Linux (CentOS 7)	GPU-System (2016)

SSH and Cluster login

- Use established *secure shell protocol (ssh)* connection to log in to the front-end cluster nodes.

  $ ssh <Username>@<Servername(or IP)>

- In the case of the RWTH Cluster:

  $ ssh <TIM>@<login18-1.hpc.itc.rwth-aachen.de>

SSH and Cluster login

- Special nodes for intensive IO operations

copy18-1.hpc.itc.rwth-aachen.de	CentOS
copy18-2.hpc.itc.rwth-aachen.de	CentOS
copy.hpc.itc.rwth-aachen.de	CentOS

- Use this nodes if you ever need to transfer big files to the cluster

Cluster – available file systems

Name	Path	Backup	Quota (file)	Quota (#files)
$HOME	*/home/*<TIM-Kennung>	yes	1000 GB	1000
$WORK	*/work/*<TIM-Kennung>	no	250 GB	1000
$HPCWORK	*/hpcwork/*<TIM-Kennung>	no	1000 GB	50000

Cluster – mount drive

- Use *sshfs* to mount remote directory to local machine:
  - $ sshfs <TIM>@copy18-1.hpc.itc.rwth-aachen.de:<Path> <local Path>
- Unmount with:
  - $ sudo umount -l <local Path>
- Only works if you are inside eduroam main-network or through RWTH VPN
  - https://doc.itc.rwth-aachen.de/display/VPN/VPN+%28Ubuntu%29+Cisco+OpenConnect

Cluster – Batch Jobs

- Since 2018 new front-end nodes, new jobs-handling system – SLURM
- SBATCH [options] command [arguments]
- Slots

Parameter	Function
-c, --cpus-per-task <numcpus>	Number of threads/processes for an OpenMP/Hybrid script
-n, --ntasks <numtasks> --ntasks-per-node <numtasks>	Number of threads/processes for an MPI job
-N, --nodes <numnodes>	Number of nodes/hosts for the job

- List of all available nodes:
  https://doc.itc.rwth-aachen.de/display/CC/Hardware+of+the+RWTH+Compute+Cluster

Cluster – Batch Jobs

- Further job parameters

Parameter	Function
-J --job-name=<jobname>	A name for the current job
--mem-per-cpu=<size>	Required RAM per allocated CPU
-o, --output=<filename>	Name for a report file, containing the standard output of the job
-t, --time=d-hh:mm:ss	Time for job execution, after this time the job is killed
-A, --account=<project>	Submit a job for a specific project
--gres=gpu:<type>:2	Requesting two GPUs per node

Cluster – Batch Jobs

- **sbatch** jobscript.sh – To run a job, specified in jobscript.sh
- At the end of the jobscript.sh file add a line for execution a script of programm
- Might need to load some modules prior, e.g.:
  - *module load python/3.6.0*
  - *module avail* – to list all available modules
- Set environment variable inside a cluster job script
  - *export PATH=/usr/local_rwth/sw/cuda/9.0.176/bin:$PATH*

Cluster – Example Scripts

```
#!/bin/bash

### Job name
#SBATCH --job-name=MYJOB

### File for the output
#SBATCH --output=MYJOB_OUTPUT

### Time your job needs to execute, e. g. 15 min 30 sec
#SBATCH --time=00:15:30

### Memory your job needs per node, e. g. 1 GB
#SBATCH --mem=1G

### The last part consists of regular shell commands:
### Change to working directory
cd /home/usr/workingdirectory

### Execute your application
myapp.exe
```

Cluster – Monitor Jobs and Resources

- Use ***squeue*** to monitor current jobs in progress:
  - squeue -u <TIM> for only jobs by specific users

- Use ***scancel*** to stop current jobs from executing
  - scancel -u <TIM> for killing jobs by specific user
  - scancel -n <Jobname> for killing a specific job

- Use ***quota*** to see occupied space on the hard-drives and used number of files.

- Use ***sinfo*** to see updated accounting information. Currently no monthly CPU-time quotas in place.

Cluster – GPU Nodes

- Login to a GPU-Node by:
  ssh <TIM>@login18-g-1.hpc.itc.rwth-aachen.de

- See GPU usage with nvidia-smi

```
Tue May  7 17:24:20 2019
+-----------------------------------------------------------------------------+
| NVIDIA-SMI 418.43       Driver Version: 418.43       CUDA Version: 10.1     |
|-------------------------------+----------------------+----------------------+
| 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 V100-SXM2...  Off  | 00000000:62:00.0 Off |                    0 |
| N/A   42C    P0    54W / 300W |      0MiB / 16130MiB |      0%   E. Process |
+-------------------------------+----------------------+----------------------+
|   1  Tesla V100-SXM2...  Off  | 00000000:89:00.0 Off |                    0 |
| N/A   42C    P0    54W / 300W |      0MiB / 16130MiB |      0%   E. Process |
+-------------------------------+----------------------+----------------------+

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

Cluster – GPU Nodes

- Login to a GPU-Node by:
  ssh <TIM>@login18-g-1.hpc.itc.rwth-aachen.de

- See GPU usage with nvidia-smi

```
Tue May  7 17:24:20 2019
+-----------------------------------------------------------------------------+
| NVIDIA-SMI 418.43       Driver Version: 418.43       CUDA Version: 10.1     |
|-------------------------------+----------------------+----------------------+
| 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 V100-SXM2...  Off  | 00000000:62:00.0 Off |                    0 |
| N/A   42C    P0    54W / 300W |      0MiB / 16130MiB |      0%   E. Process |
+-------------------------------+----------------------+----------------------+
|   1  Tesla V100-SXM2...  Off  | 00000000:89:00.0 Off |                    0 |
| N/A   42C    P0    54W / 300W |      0MiB / 16130MiB |      0%   E. Process |
+-------------------------------+----------------------+----------------------+

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

GPU-id and type

Cluster – GPU Nodes

- Login to a GPU-Node by:
  ssh <TIM>@login18-g-1.hpc.itc.rwth-aachen.de

- See GPU usage with nvidia-smi

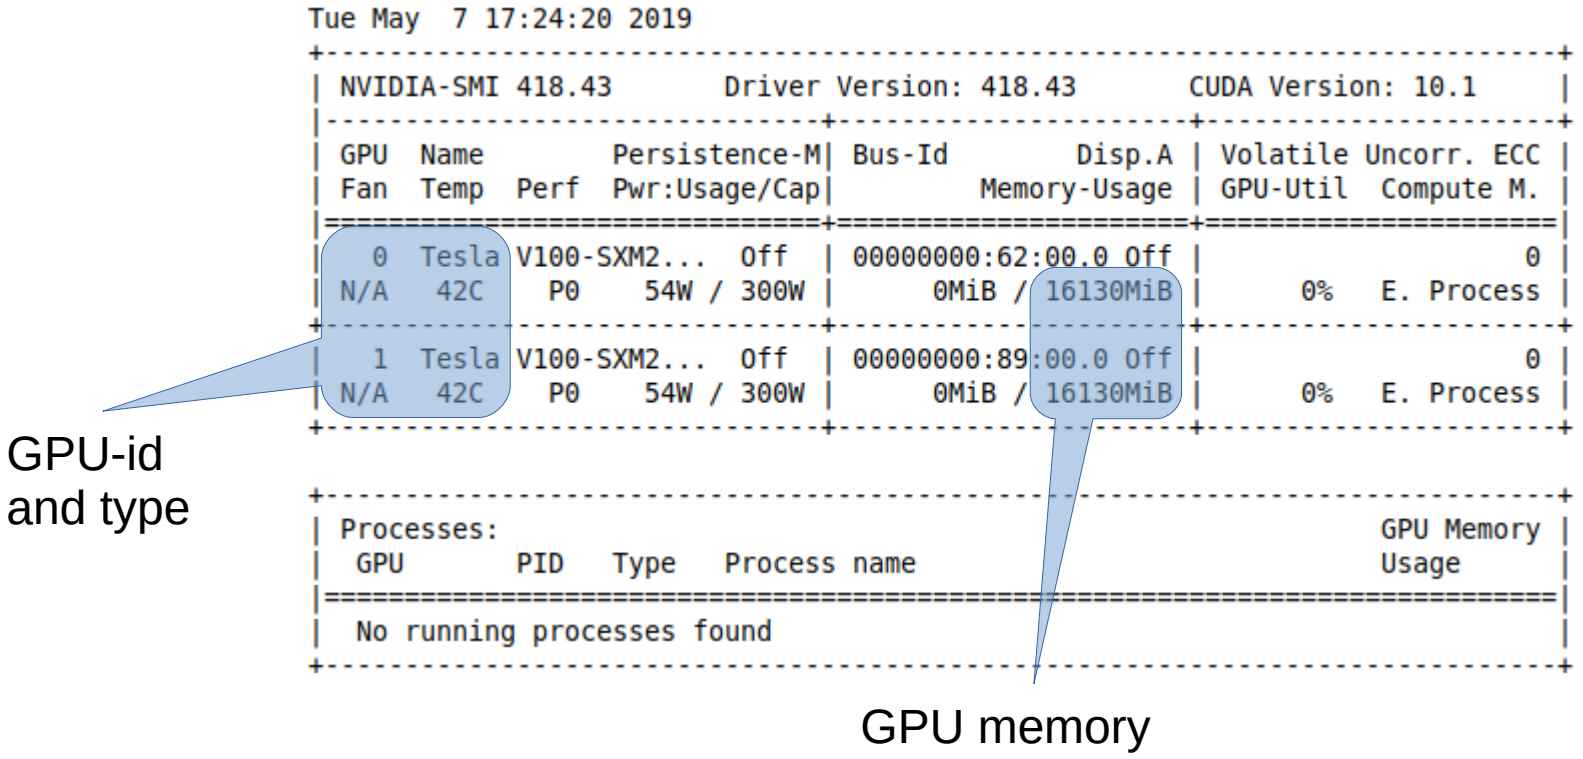

Cluster – GPU Nodes

- Login to a GPU-Node by:
  ssh <TIM>@login18-g-1.hpc.itc.rwth-aachen.de

- See GPU usage with nvidia-smi

Compute model: 1 person

```
Tue May  7 17:24:20 2019
+-----------------------------------------------------------------------------+
| NVIDIA-SMI 418.43       Driver Version: 418.43       CUDA Version: 10.1     |
|-------------------------------+----------------------+----------------------+
| 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 V100-SXM2...  Off  | 00000000:62:00.0 Off |                    0 |
| N/A   42C    P0    54W / 300W |      0MiB / 16130MiB |      0%   E. Process |
+-------------------------------+----------------------+----------------------+
|   1  Tesla V100-SXM2...  Off  | 00000000:89:00.0 Off |                    0 |
| N/A   42C    P0    54W / 300W |      0MiB / 16130MiB |      0%   E. Process |
+-------------------------------+----------------------+----------------------+

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

GPU-id and type

GPU memory

Cluster – Exercise – GPU with python

- Sign in to the cluster
  $ ssh <TIM>@login18-g-2.hpc.itc.rwth-aachen.de

- Check if you're successfully in the group:
  $ groups # nova0028 should be there

- Load python3.6 module and install some libraries locally
  $ module load python/3.6.0
  $ pip3.6 install --user matplotlib tensorflow_gpu keras

- Create a new directory and copy the files needed
  $ mkdir BioInfo
  $ rsync -r /hpcwork/nova0028/BioInfoLab2019/examples BioInfo
  $ cd BioInfo/examples/GPU

- Make start_job.sh executable ($ chmod 770 <filename>)
- **Edit** (with **vim**) your email address and submit job ($ sbatch start_job.sh)
- Check status regularly ($ squeue -u <TIM>)

Cluster – Exercise – CPU with R

- Open R and install some packages
  $ R
  > install.packages(“dplyr”)
  > install.packages(“Seurat’)

- Navigate to the folder
  $ cd ~/BioInfo/examples/CPU

- Make start_job.sh executable ($ chmod 770 <filename>)
- **Edit** (with vim) your email address and submit job ($ sbatch start_job.sh)
- Check status regularly ($ squeue -u <TIM>)

Further Readings

- Linux tutorials:
  - https://www.tutorialspoint.com/unix/
  - https://ryanstutorials.net/linuxtutorial/

- SLURM – batch system
  - https://slurm.schedmd.com/documentation.html
  - https://doc.itc.rwth-aachen.de/display/CC/Using+the+SLURM+Batch+System
  - https://hpc-wiki.info/hpc/SLURM#Jobscript_Examples