

High Performance Computing

ADVANCED SCIENTIFIC COMPUTING

Prof. Dr. – Ing. Morris Riedel

Adjunct Associated Professor

School of Engineering and Natural Sciences, University of Iceland, Reykjavik, Iceland

Research Group Leader, Juelich Supercomputing Centre, Forschungszentrum Juelich, Germany

PRACTICAL LECTURE 0.1

[in @Morris Riedel](#)

[@MorrisRiedel](#)

[@MorrisRiedel](#)

Short Introduction to UNIX & SSH

August 29, 2019

Webinar



UNIVERSITY OF ICELAND
SCHOOL OF ENGINEERING AND NATURAL SCIENCES
FACULTY OF INDUSTRIAL ENGINEERING,
MECHANICAL ENGINEERING AND COMPUTER SCIENCE



JÜLICH
Forschungszentrum

JÜLICH
SUPERCOMPUTING
CENTRE




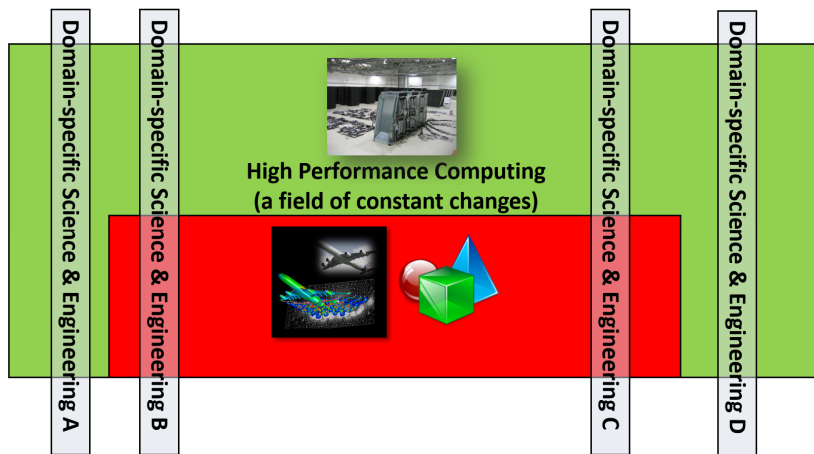
HELMHOLTZ
RESEARCH FOR GRAND CHALLENGES






HELMHOLTZ
ARTIFICIAL INTELLIGENCE
COOPERATION UNIT

Review of Lecture 0 – Prologue

■ Course Motivation & Information



■ Course Organization & Content



[4] Distributed & Cloud Computing Book

Outline of the Course

1. High Performance Computing
2. Parallel Programming with MPI
3. Parallelization Fundamentals
4. Advanced MPI Techniques
5. Parallel Algorithms & Data Structures
6. Parallel Programming with OpenMP
7. Graphical Processing Units (GPUs)
8. Parallel & Scalable Machine & Deep Learning
9. Debugging & Profiling & Performance Toolsets
10. Hybrid Programming & Patterns

11. Scientific Visualization & Scalable Infrastructures
12. Terrestrial Systems & Climate
13. Systems Biology & Bioinformatics
14. Molecular Systems & Libraries
15. Computational Fluid Dynamics & Finite Elements
16. Epilogue

+ additional practical lectures & Webinars for our hands-on assignments in context

- Practical Topics
- Theoretical / Conceptual Topics

Outline

■ Using UNIX on HPC Systems

- HPC Systems & Modular Supercomputing Architecture
- HPC System Examples DEEP & Jötunn
- Selected UNIX Commands
- Module Environment
- Basic Editor VI


■ Using SSH to connect to HPC Systems

- SSH Clients
- Private/Public Key Pairs
- SSH Key Example Login HPC System DEEP
- Username & Password Example HPC System Jötunn
- Workaround for external login to HPC System Jötunn

- This lecture is not considered to be a full introduction to UNIX and SSH and rather focusses on selected commands and concepts relevant for assignments
- The goal of this lecture is to make course participants aware of the UNIX environments existing on world-wide HPC systems and how to connect to them using the SSH protocol



Using UNIX on HPC Systems




Selected Learning Outcomes – Revisited

- Students understand...
 - Latest developments in **parallel processing** & **high performance computing (HPC)**
 - **How to create and use high-performance clusters**
 - What are **scalable networks** & **data-intensive workloads**
 - The importance of **domain decomposition**
 - **Complex aspects of parallel programming**
 - **HPC environment tools** that support programming or analyze behaviour
 - Different abstractions of **parallel computing on various levels**
 - Foundations and approaches of **scientific domain-specific applications**
- Students are able to ...
 - Programm and use HPC programming paradigms
 - Take advantage of innovative scientific computing simulations & technology
 - Work with technologies and tools to handle parallelism complexity



Understanding HPC Systems – Revisited (cf. Lecture 0 Prologue)

- High Performance Computing (HPC) is based on computing resources that enable the efficient use of parallel computing techniques through specific support with dedicated hardware such as high performance cpu/core interconnections.



- High Throughput Computing (HTC) is based on commonly available computing resources such as commodity PCs and small clusters that enable the execution of 'farming jobs' without providing a high performance interconnection between the cpu/cores.



➤ The complementary Cloud Computing & Big Data – Parallel Machine & Deep Learning Course focusses on High Throughput Computing

HPC & Data-intensive Sciences – A Field of Constant Evolution – Revisited

1.000.000 FLOP/s

~1984




© Photograph by Rama,
Wikimedia Commons

- Floating Point Operations per one second (FLOPS or FLOP/s)
- 1 GigaFlop/s = 10^9 FLOPS
- 1 TeraFlop/s = 10^{12} FLOPS
- 1 PetaFlop/s = 10^{15} FLOPS
- 1 ExaFlop/s = 10^{18} FLOPS

1.000.000.000.000.000 FLOP/s

~295.000 cores ~2009 (JUGENE)



Upgrade JUGENE to JUQUEEN



>5.900.000.000.000.000 FLOP/s

~ 500.000 cores

~2013 → end of service in 2018

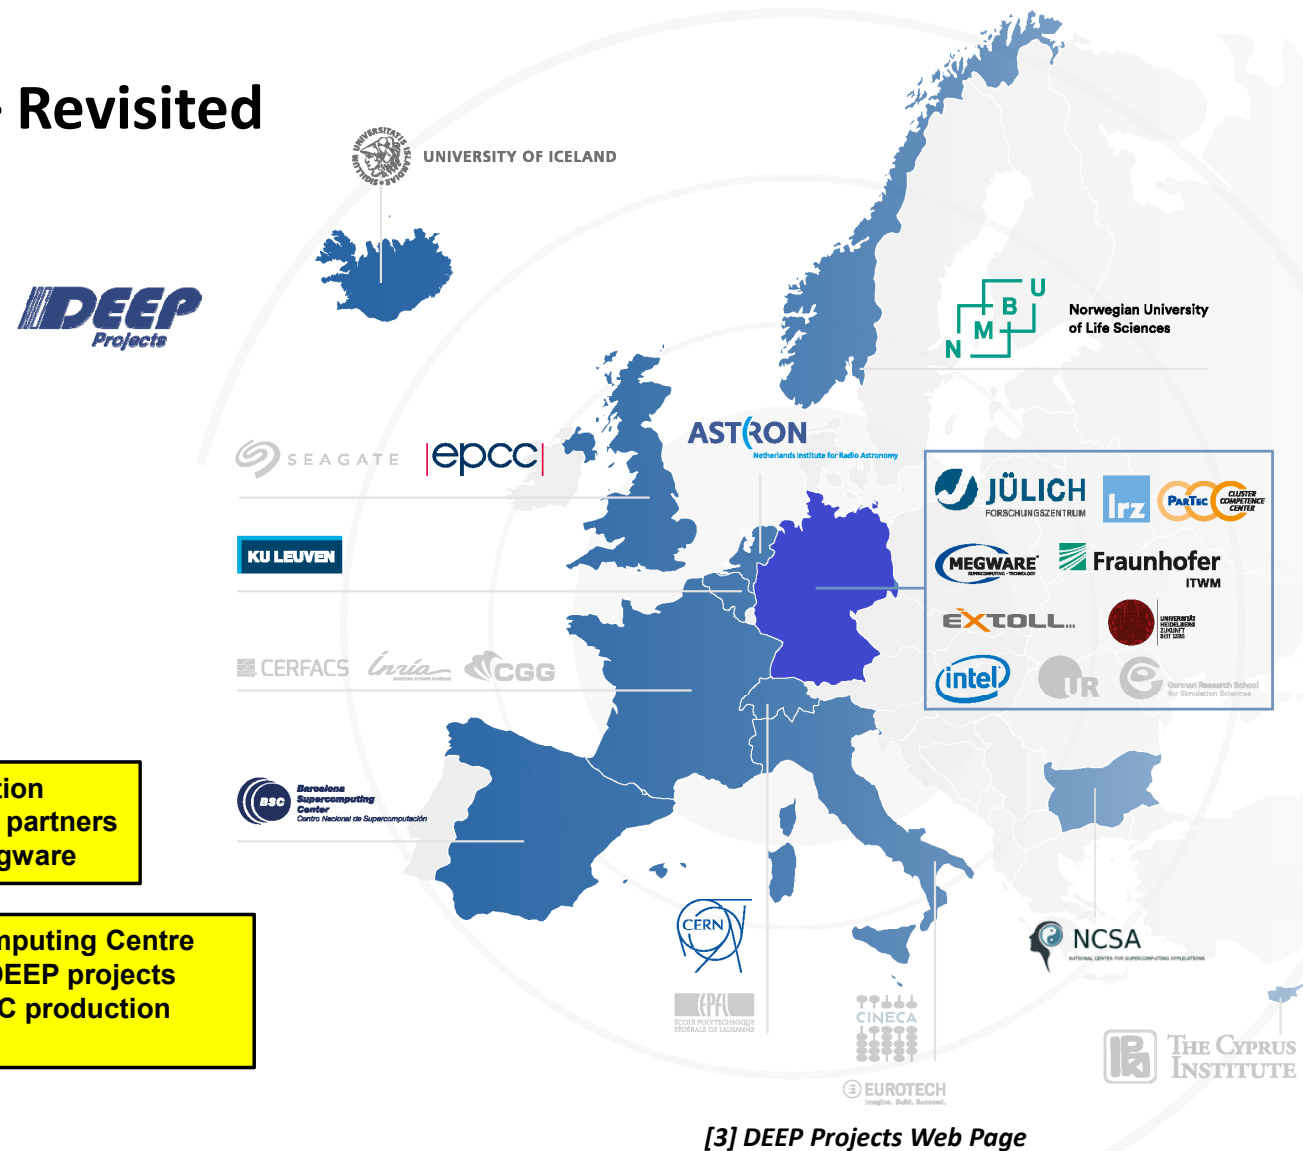
DEEP series of PROJECTS & HPC – Revisited



- 3 EU Exascale projects
DEEP, DEEP-ER, DEEP-EST
- 27 partners
Coordinated by JSC
- EU-funding: 30 M€
JSC-part > 5,3 M€
- Nov 2011 – Dec 2020

- **Strong collaboration with our industry partners Intel, Extoll & Megware**

- Juelich Supercomputing Centre implements the DEEP projects designs in its HPC production infrastructure





IBM

IBM Power 6
JUMP, 9 TFlop/s

JUROPA
200 TFlop/s
HPC-FF
100 TFlop/s

JURECA Cluster (2015)
2.2 PFlop/s



IBM Power 4+
JUMP (2004), 9 TFlop/s



IBM Blue Gene/L
JUBL, 45 TFlop/s



IBM Blue Gene/P
JUGENE, 1 PFlop/s



IBM Blue Gene/Q
JUQUEEN (2012)
5.9 PFlop/s



File
Server
GPFS,



Proof of
Concept in European
DEEP Project




JURECA Booster (2017)
5 PFlop/s

NVIDIA

JUWELS_Cluster
Module (2018)
12 PFlop/s



Hierarchical
Storage Server
Modular
Supercomputer



JUWELS_Scalable
Module (2019/20)
50+ PFlop/s

General Purpose Cluster

Highly scalable

Atos
BULL

PARTEC CLUSTER
COMPETENCE
CENTER

HPC System – DEEP Testcluster

■ DEEP-EST Project Prototype HPC System


- Dynamical Exascale Entry Platform (DEEP)

■ Implements Modular Supercomputing Architecture (MSA)

- Cluster module (installed at JSC in April 2019)
- Extreme Scale Booster module (to be installed during 2019)
- Data Analytics module (to be installed during 2019, prototype available)

■ Data Analytics Module (DAM) prototype

- 3 nodes x 4 Graphics Processing Units (GPUs)
- Each node with 4 x Nvidia Tesla V100 GPU
- Each node with host CPUs:
2 x Intel Xeon 'Skylake' Silver 4112
- Access via common DEEP environment



[4] Distributed & Cloud Computing Book




[2] DEEP Test Cluster

➤ We will use the modular supercomputing architecture on the DEEP test cluster with deep learning and Graphic Processing Units (GPUs)

HPC System – Jötunn Cluster


- 4 Nodes
 - Cpu: 2x Intel Xeon CPU E5-2690 v3 @ 2.60GHz (2.6 GHz, 12 core)
- Memory
 - 128GB DDR4
- Interconnect
 - 10 Gb/s Ethernet
- Ganglia monitoring service
 - Shows usage of CPUs



Jötunn

Jötunn at a glance

nodes: 4
cpu: 2x Intel Xeon CPU E5-2690 v3 @ 2.60GHz (2.6 GHz, 12 core)
memory: 128GB DDR4
Interconnect: 10 Gb/s ethernet



[1] Icelandic HPC Machines & Community

➤ We will have a visit to computing room of Jötunn to 'touch metal' and will meet our HPC System expert Hjörleifur Sveinbjörnsson

Selected UNIX Commands: hostname

- Example: 'hostname -A' command on our Jötunn cluster

```
[morris@jotunn ~]$ hostname -A  
jotunn-login2.rhi.hi.is jotunn jotunn.rhi.hi.is  
[morris@jotunn ~]$ exit  
logout  
Connection to jotunn.rhi.hi.is closed.
```

- Example: 'hostname -A' command on our DEEP Testcluster

```
[riedell@deepv ~]$ hostname -A  
deepv.deep zam1015.zam.kfa-juelich.de deepvm.deep
```



[2] DEEP Test Cluster



[1] Icelandic HPC Machines & Community

Selected UNIX Commands: whoami & clear

- Example: 'whoami' command on our DEEP Testcluster

```
[riedell@deepv ~]$ whoami  
riedell
```



- Example: 'clear' command on our DEEP Testcluster

```
Last login: Thu Aug 29 08:20:33 2019 from pd95b302c.dip0.t-ipconnect.de  
#####  
# DEEP Testcluster  
#####  
# Useful job status tools:  
#  
# Overall node usage: slurmtop  
# Show jobs in the system: squeue  
# List reservations: scontrol show res  
# Check job issues: scontrol show job <jobid>  
# See /etc/slurm/README for details and known problems  
#  
#####  
# /usr/local now served by local beegfs file system.  
# performance issues under investigation.  
# pn Thu Jun 13 14:26:48 CEST 2019  
#  
#####  
# The transition to the new software stack will happen on Tuesday 9th  
# To enable it now use "enable_new_software_stack"  
# To enable the current stack use "enable_old_stack"  
# To enable the legacy stack use "enable_legacy_stack"  
#  
#####  
# /usr/local on deepv, dp-cn nodes:  
# now mounted without acls and extended attributes.  
# cm/pn Fri Jul 26 10:30:50 CEST 2019  
#  
#####  
[riedell@deepv ~]$ whoami  
riedell  
[riedell@deepv ~]$ hostname -f  
deepv.deep  
[riedell@deepv ~]$ hostname -A  
deepv.deep zam1015.zam.kfa-juelich.de deepvm.deep  
[riedell@deepv ~]$ clear
```



[1] Icelandic HPC Machines & Community

Selected UNIX Commands: ls & pwd

- Example: 'ls -al' command on our DEEP Testcluster

```
[riedell@deepv ~]$ ls -al
total 14
drwx----- 10 riedell jusers 4096 Aug 29 08:20 .
drwxr-xr-x  8 root    root   4096 May 14 17:03 ..
-rw-----  1 riedell jusers 2613 Aug 16 17:40 .bash_history
-rwx-----  1 riedell jusers  535 Nov 30  2018 .bash_profile
-rwx-----  1 riedell jusers  501 Nov 30  2018 .bashrc
drwx-----  6 riedell jusers 4096 Aug 15 13:30 .cache
drwxr-xr-x  3 riedell jusers 4096 Aug 15 13:30 .config
drwxr-xr-x  6 riedell jusers 4096 Aug 15 19:51 hugi
drwxr-xr-x  2 riedell jusers 4096 Aug 15 13:30 .keras
drwxr-xr-x  3 riedell jusers 4096 Aug 15 11:54 .lmod.d
drwx-----  4 riedell jusers 4096 Aug 15 13:20 .local
drwxr-xr-x  5 riedell jusers 4096 Aug 15 14:00 project
lrwxrwxrwx  1 riedell jusers   9 Nov 30  2018 shared -> ../shared
drwx-----  2 riedell jusers 4096 Nov 30  2018 .ssh
-rw-----  1 riedell jusers 2530 Aug 15 14:00 .viminfo
-rw-----  1 riedell jusers  204 Aug 29 08:20 .Xauthority
```



- Example: 'pwd' command on our DEEP Testcluster

```
[riedell@deepv ~]$ cd HPC-Course-Fall-2019/
[riedell@deepv HPC-Course-Fall-2019]$ pwd
/p/home/jusers/riedell/deep/HPC-Course-Fall-2019
```



[2] DEEP Test Cluster

Selected UNIX Commands: mkdir & cd

- Example: 'mkdir FOLDER' and 'cd FOLDER' command on our DEEP Testcluster

```
[riedell@deepv ~]$ mkdir HPC-Course-Fall-2019
[riedell@deepv ~]$ ls -al
total 15
drwx----- 11 riedell jusers 4096 Aug 29 09:30 .
drwxr-xr-x  8 root    root    4096 May 14 17:03 ..
-rw-----  1 riedell jusers 2713 Aug 29 09:30 .bash_history
-rwx-----  1 riedell jusers  535 Nov 30 2018 .bash_profile
-rwx-----  1 riedell jusers  501 Nov 30 2018 .bashrc
drwx-----  6 riedell jusers 4096 Aug 15 13:30 .cache
drwxr-xr-x  3 riedell jusers 4096 Aug 15 13:30 .config
drwxr-xr-x  2 riedell jusers 4096 Aug 29 09:30 HPC-Course-Fall-2019
drwxr-xr-x  6 riedell jusers 4096 Aug 15 19:51 hugi
drwxr-xr-x  2 riedell jusers 4096 Aug 15 13:30 .keras
drwxr-xr-x  3 riedell jusers 4096 Aug 15 11:54 .lmod.d
drwx-----  4 riedell jusers 4096 Aug 15 13:20 .local
drwxr-xr-x  5 riedell jusers 4096 Aug 15 14:00 project
lrwxrwxrwx  1 riedell jusers   9 Nov 30 2018 shared -> ../shared
drwx-----  2 riedell jusers 4096 Aug 29 09:28 .ssh
-rw-----  1 riedell jusers 2530 Aug 15 14:00 .viminfo
-rw-----  1 riedell jusers  204 Aug 29 09:30 .Xauthority
[riedell@deepv ~]$ cd HPC-Course-Fall-2019
[riedell@deepv HPC-Course-Fall-2019]$ ls -al
total 2
drwxr-xr-x  2 riedell jusers 4096 Aug 29 09:30 .
drwx----- 11 riedell jusers 4096 Aug 29 09:30 ..
```



[2] DEEP Test Cluster

HPC System Module Environment

- Knowledge of **installed compilers** essential (e.g. C, Fortran90, etc.)
 - Different versions and types of compilers exist (Intel, GNU, MPI, etc.)
 - E.g. **mpicc** pingpong.c -o pingpong
- **Module** environment tool
 - Avoids to manually setup environment information for every application
 - Simplifies shell initialization and lets users easily modify their environment
 - Modules can be loaded and unloaded
 - Enable the installation of software in different versions
- **Module avail**
 - Lists all available modules on the HPC system (e.g. compilers, MPI, etc.)
- **Module load**
 - Loads particular modules into the current work environment
 - E.g. module load gnu openmpi



[2] DEEP Test Cluster

HPC System Module Environment: module avail & module load

- Example: 'module avail' & 'module load' command on our DEEP Testcluster

```
[riedell@deepv ~]$ module avail

----- /usr/local/software/skylake/Stages/2019a/modules/all/Compiler/GCCcore/8.3.0 -----
Autotools/20180311      PAPI/5.7.0
Bazel/0.20.0            Perl/5.28.1
CMake/3.14.0            PostgreSQL/11.2
Clang/8.0.0-GCC-8.3.0-CUDA-10.1.105 (g)  Python/2.7.16
CubeGUI/4.4.3           Python/3.6.8 (D)
CubeLib/4.4.3           SciPy-Stack/2019a-Python-2.7.16
CubeWriter/4.4.2        SciPy-Stack/2019a-Python-3.6.8 (D)
Doxygen/1.8.15          Tcl/8.6.9
GEOS/3.7.1-Python-3.6.8 TensorFlow/1.13.1-GPU-Python-3.6.8 (g)
GMP/6.1.2              UCX/1.5.1
Graphviz/2.40.1         X11/20190311
HDF5/1.10.5-serial      cURL/7.64.1
LLVM/8.0.0             flex/2.6.4
Mesa/19.0.1            h5py/2.9.0-serial-Python-3.6.8
Meson/0.50.0-Python-3.6.8 numactl/2.0.12
NCCL/2.4.6-1-CUDA-10.1.105 (g)  unzip/6.0
Ninja/1.9.0

----- /usr/local/software/skylake/Stages/2019a/UI/Compilers -----
GCC/8.3.0      Intel/2019.3.199-GCC-8.3.0

----- /usr/local/software/skylake/Stages/2019a/UI/Tools -----
Advisor/2019_update3  Inspector/2019_update3  cuDNN/7.5.1.10-CUDA-10.1.105 (g)
CUDA/10.1.105 (g)     JUBE/2.2.2             intel-para/2019a-mt
EasyBuild/3.8.1       Java/1.8               intel-para/2019a (D)
EasyBuild/3.9.1 (D)   VTune/2019_update3     tbb/2019.4.199

----- /usr/local/software/extoll -----
extoll  openmpi/1.6.1_extoll

Where:
g: built for GPU
D: Default Module

Use "module spider" to find all possible modules.
Use "module keyword key1 key2 ..." to search for all possible modules matching any of the "keys".

[riedell@deepv HPC-Course-Fall-2019]$ module load Python/3.6.8
```



[2] DEEP Test Cluster

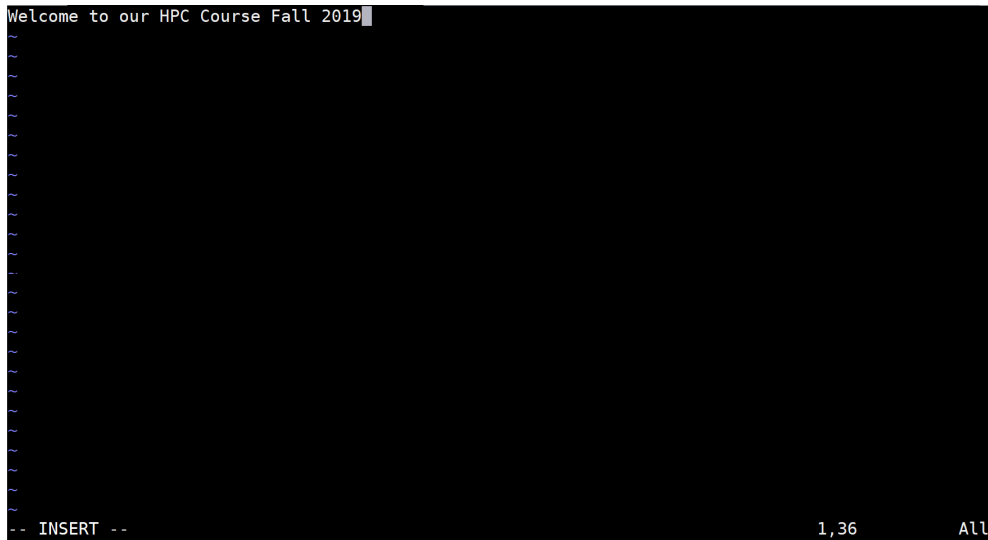
HPC System Environment Basic Editor VI

- HPC Systems have often not a GUI editor
 - Simple editor VI is always available
- Example: `'vi FILENAME'` using command `':a'` to insert

```
[riedell@deepv HPC-Course-Fall-2019]$ vi testfile
```



```
Welcome to our HPC Course Fall 2019
```


A screenshot of the vi editor interface. The top line shows the message "Welcome to our HPC Course Fall 2019". The rest of the screen is black. At the bottom left, it says "-- INSERT --". At the bottom right, it shows the cursor position "1,36" and the file name "All".

```
-- INSERT -- 1,36 All
```



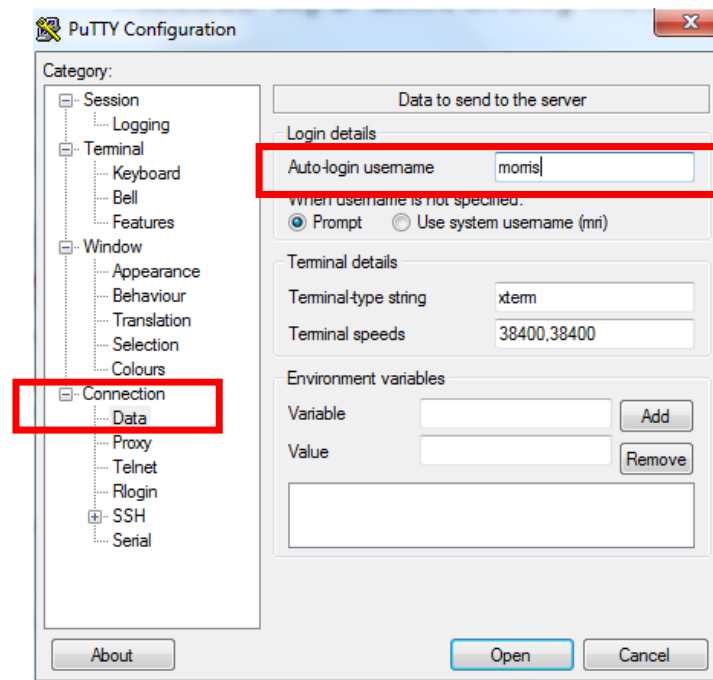
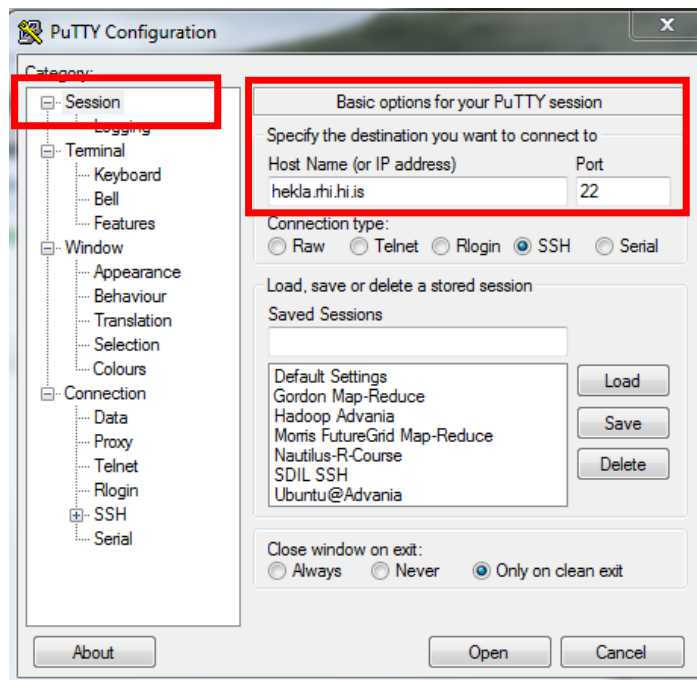
[2] DEEP Test Cluster

Using SSH to connect to HPC Systems

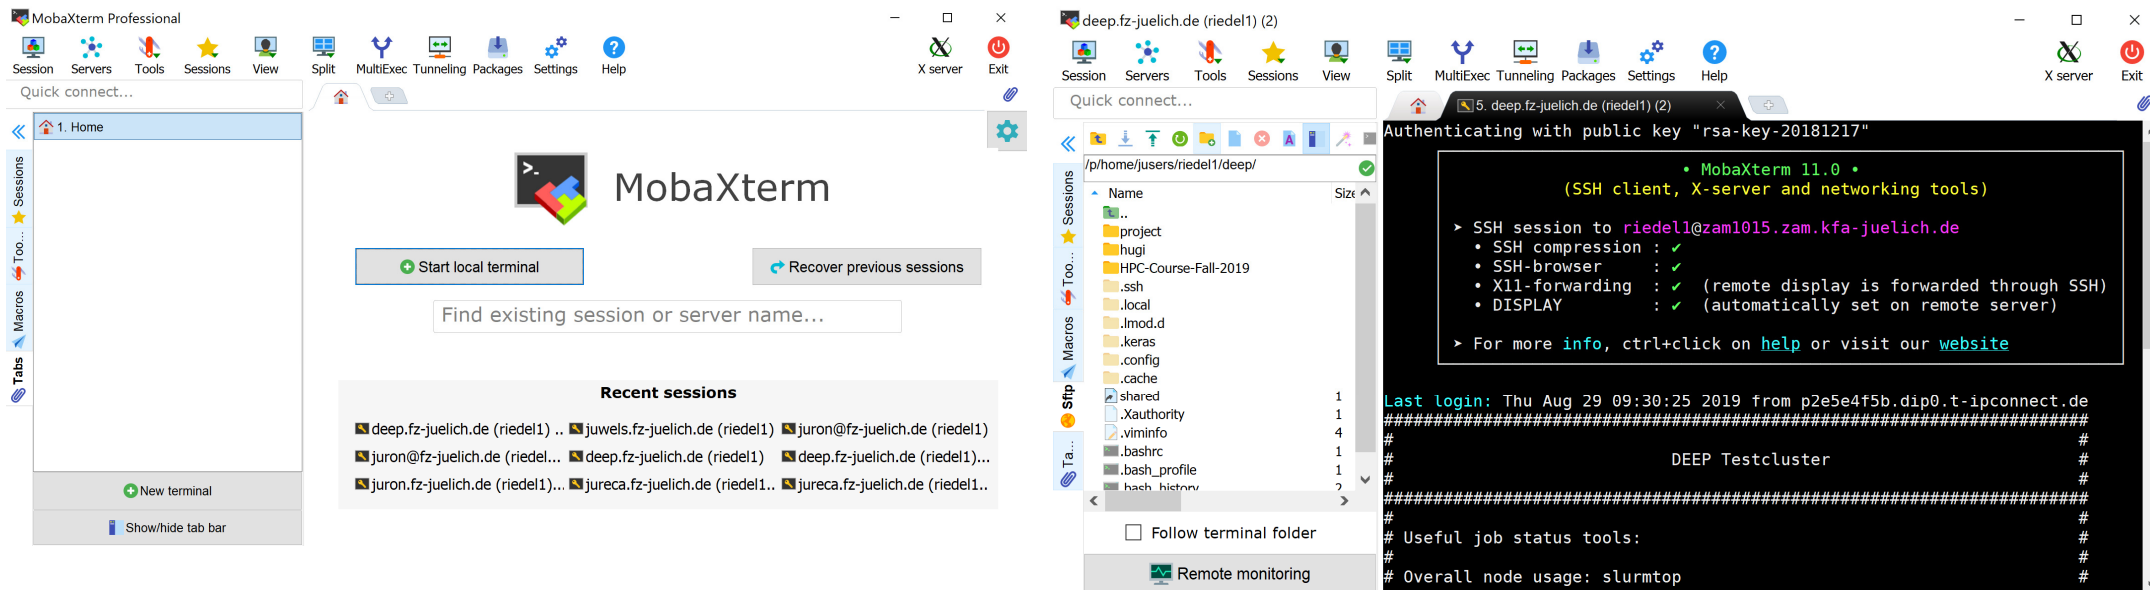


SSH Clients – Putty for Windows

- Example: Putty SSH Client for Windows
 - Not recommended, better install MobaXterm



MobaXterm SSH Client



[5] MobaXterm SSH Client

SSH Keys – Example Morris Public Key

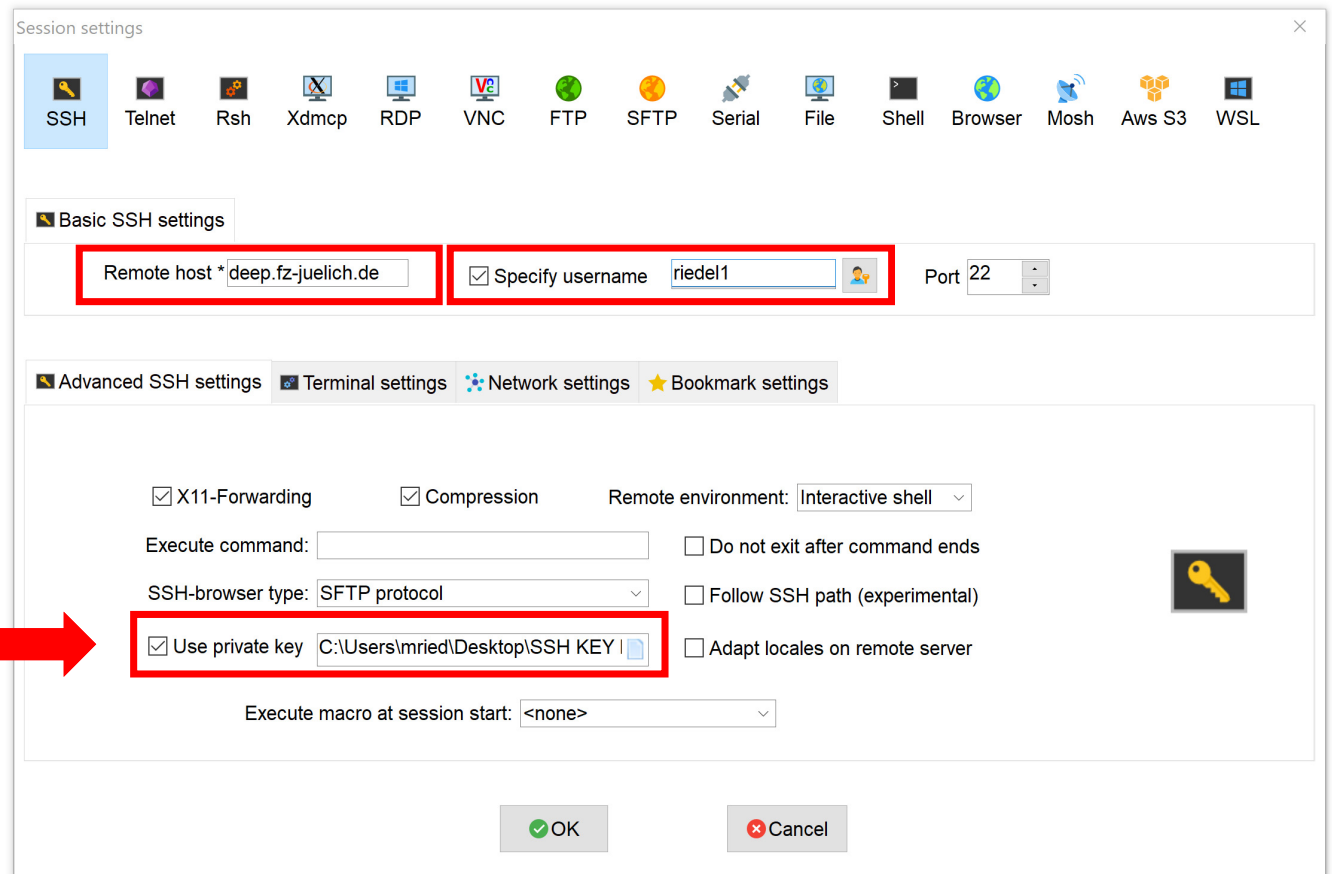
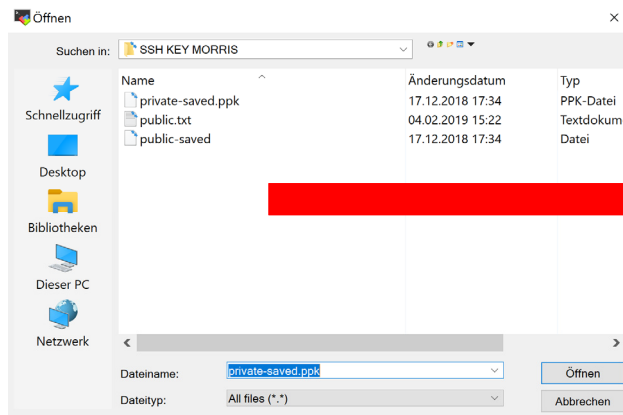
- Everybody is allowed to see the **public key**
 - Given to HPC administrators
 - Sometimes uploaded in Web forms
 - **Only Morris has the corresponding private key that never leaves the laptop of Morris (!)**

- SSH Key Example

- ssh-rsa
AAAAB3NzaC1yc2EAAAABJQAAAQEArU2IJQmVEwVjsQ6N9PUJP0KukCGQV2yAMs3hop0stsvfb4Iac7s2Pqk
wOgoFPZGwRCSGcA2/rISJX3MxEmx7EQLD5sw63r8LqvETXy4hmeflIBwpcIxMBYSLujWdCH9K60Q6TApMz4h
V+fsZRIgbTx7hs9Y2a3TiiSE032lvzxMYTvw8NYlhXOP9PzTR1jebVj3rgcOIYLPmGzl4YIbCZJVleJlwfkZscOH9zT4
KI5SpQuk5Q+LyMI95X3xsk3xPMCuocqsYmIY6Gp+BCAYJsdCXFNDJ3SCcphziTqrE+F2EroI4AoegVIH/vhPaAZg
Q222nV2rDsN+uDhaBf+76Q== rsa-key-20181217

SSH Keys – Use Private/Public Key Pair to Access DEEP HPC System

- Remember to use your **Private SSH Key** to connect to the DEEP system
 - Corresponding **Public SSH key** is already uploaded on the HPC System (remote host) per username(!)



SSH Access to HPC Systems – DEEP HPC System Example

```
Authenticating with public key "rsa-key-20181217"

      • MobaXterm 11.0 •
      (SSH client, X-server and networking tools)

> SSH session to riedell@zam1015.zam.kfa-juelich.de
• SSH compression : ✓
• SSH-browser      : ✓
• X11-forwarding   : ✓ (remote display is forwarded through SSH)
• DISPLAY          : ✓ (automatically set on remote server)

> For more info, ctrl+click on help or visit our website

Last login: Fri Aug 16 16:36:51 2019 from zam106.zam.kfa-juelich.de
#####
#                               #
#          DEEP Testcluster     #
#                               #
#####
# Useful job status tools:      #
#                               #
# Overall node usage: slurmtop  #
# Show jobs in the system: queue #
# List reservations: scontrol show res #
# Check job issues: scontrol show job <jobid> #
# See /etc/slurm/README for details and known problems #
#                               #
#####
# /usr/local now served by local beegfs file system. #
# performance issues under investigation.            #
# pn Thu Jun 13 14:26:48 CEST 2019                   #
#                               #
#####
# The transition to the new software stack will happen on Tuesday 9th #
# To enable it now use "enable_new_software_stack"                  #
# To enable the current stack use "enable_old_stack"                 #
# To enable the legacy stack use "enable_legacy_stack"               #
#                               #
#####
# /usr/local on deepv, dp-cn nodes:                               #
# now mounted without acls and extended attributes.                #
# cm/pn Fri Jul 26 10:30:50 CEST 2019                               #
#                               #
#####
[riedell@deepv ~]$
```

- HPC System Address
 - deep.fz-juelich.de
 - alias for zam1015.zam.kfa-juelich.de
- HPC System Username
 - Example: [riedell](#)
 - Every student will get a different username
- HPC System Welcome Screen
 - If SSH login was succesful
 - Shows useful information about the system
 - E.g. status of the file system or known errors / bugs
 - E.g. help with important commands

SSH Access to HPC System – Jötunn HPC System Example (1)

- Nodes


- 4 cpu: 2x Intel Xeon CPU E5-2690 v3 @ 2.60GHz (2.6 GHz, 12 core)

- Memory

- 128GB DDR4

- Interconnect

- 10 Gb/s ethernet



- Access via accounts (accounts planned to be ready next week)

- ssh username@jotunn.rhi.hi.is
- Only reachable within University network
- From outside use first ssh uglusername@hekla.rhi.hi.is (UGLA account), then ssh username@jotunn.rhi.hi.is

[1] Icelandic HPC Machines & Community

➤ We will have a visit to computing room of Jötunn to 'touch metal' and will meet our HPC System expert Hjörleifur Sveinbjörnsson

SSH Access to HPC System – Jötunn HPC System Example (1)

■ Example: first login via Hekla

```
[morris@hekla ~]$ ssh morris@hekla.rhi.hi.is
The authenticity of host 'hekla.rhi.hi.is (2a00:c88:4000:1650::165:2)' can't be established.
RSA key fingerprint is 03:d4:9c:06:7e:0e:56:f4:aa:e3:f0:fe:57:bb:e7:12.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'hekla.rhi.hi.is,2a00:c88:4000:1650::165:2' (RSA) to the list of known hosts.

-----
Thu ert ad tengjast Heklu (hekla.rhi.hi.is) ffolnotendavel RHI.
Fyrir alla nemendur og starfsmenn Haskola Islands.
Leidbeiningar: http://rhi.hi.is/fjolnotendatolur

You are connecting Hekla (hekla.rhi.hi.is) for all students and
staff of the University of Iceland.
Instructions: http://rhi.hi.is/multi_user_computers
-----

morris@hekla.rhi.hi.is's password:
Last login: Tue Sep  5 08:50:28 2017 from 109.133.53.203

Styrikerfi: GNU/Linux
CentOS release 6.8 (Final)

Fjöldi tengdra notenda: 3
[morris@hekla ~]$
```



```
[morris@hekla ~]$ ssh morris@jotunn.rhi.hi.is
morris@jotunn.rhi.hi.is's password:
Last login: Tue Sep  5 04:10:01 2017 from hekla.rhi.hi.is
Welcome to Jötunn

See the jotunn sections at http://ihpc.is


Each user has 100G quota so be tidy and
back up your files

[morris@jotunn ~]$
```



[1] Icelandic HPC Machines & Community

Lecture Bibliography



Lecture Bibliography

- [1] Icelandic HPC Machines & Community, Online:
<http://ihpc.is>
- [2] DEEP-EST Project DEEP Test Cluster, Online:
<https://www.fz-juelich.de/ias/jsc/EN/Expertise/Supercomputers/DEEP-EST/node.html>
- [3] DEEP Projects Web page, Online:
<http://www.deep-projects.eu/>
- [4] K. Hwang, G. C. Fox, J. J. Dongarra, 'Distributed and Cloud Computing', Book, Online:
http://store.elsevier.com/product.jsp?locale=en_EU&isbn=9780128002049
- [5] MobaXterm SSH Client, Online:
<https://mobaxterm.mobatek.net/>

