

Quick start using a CNCA supercomputer

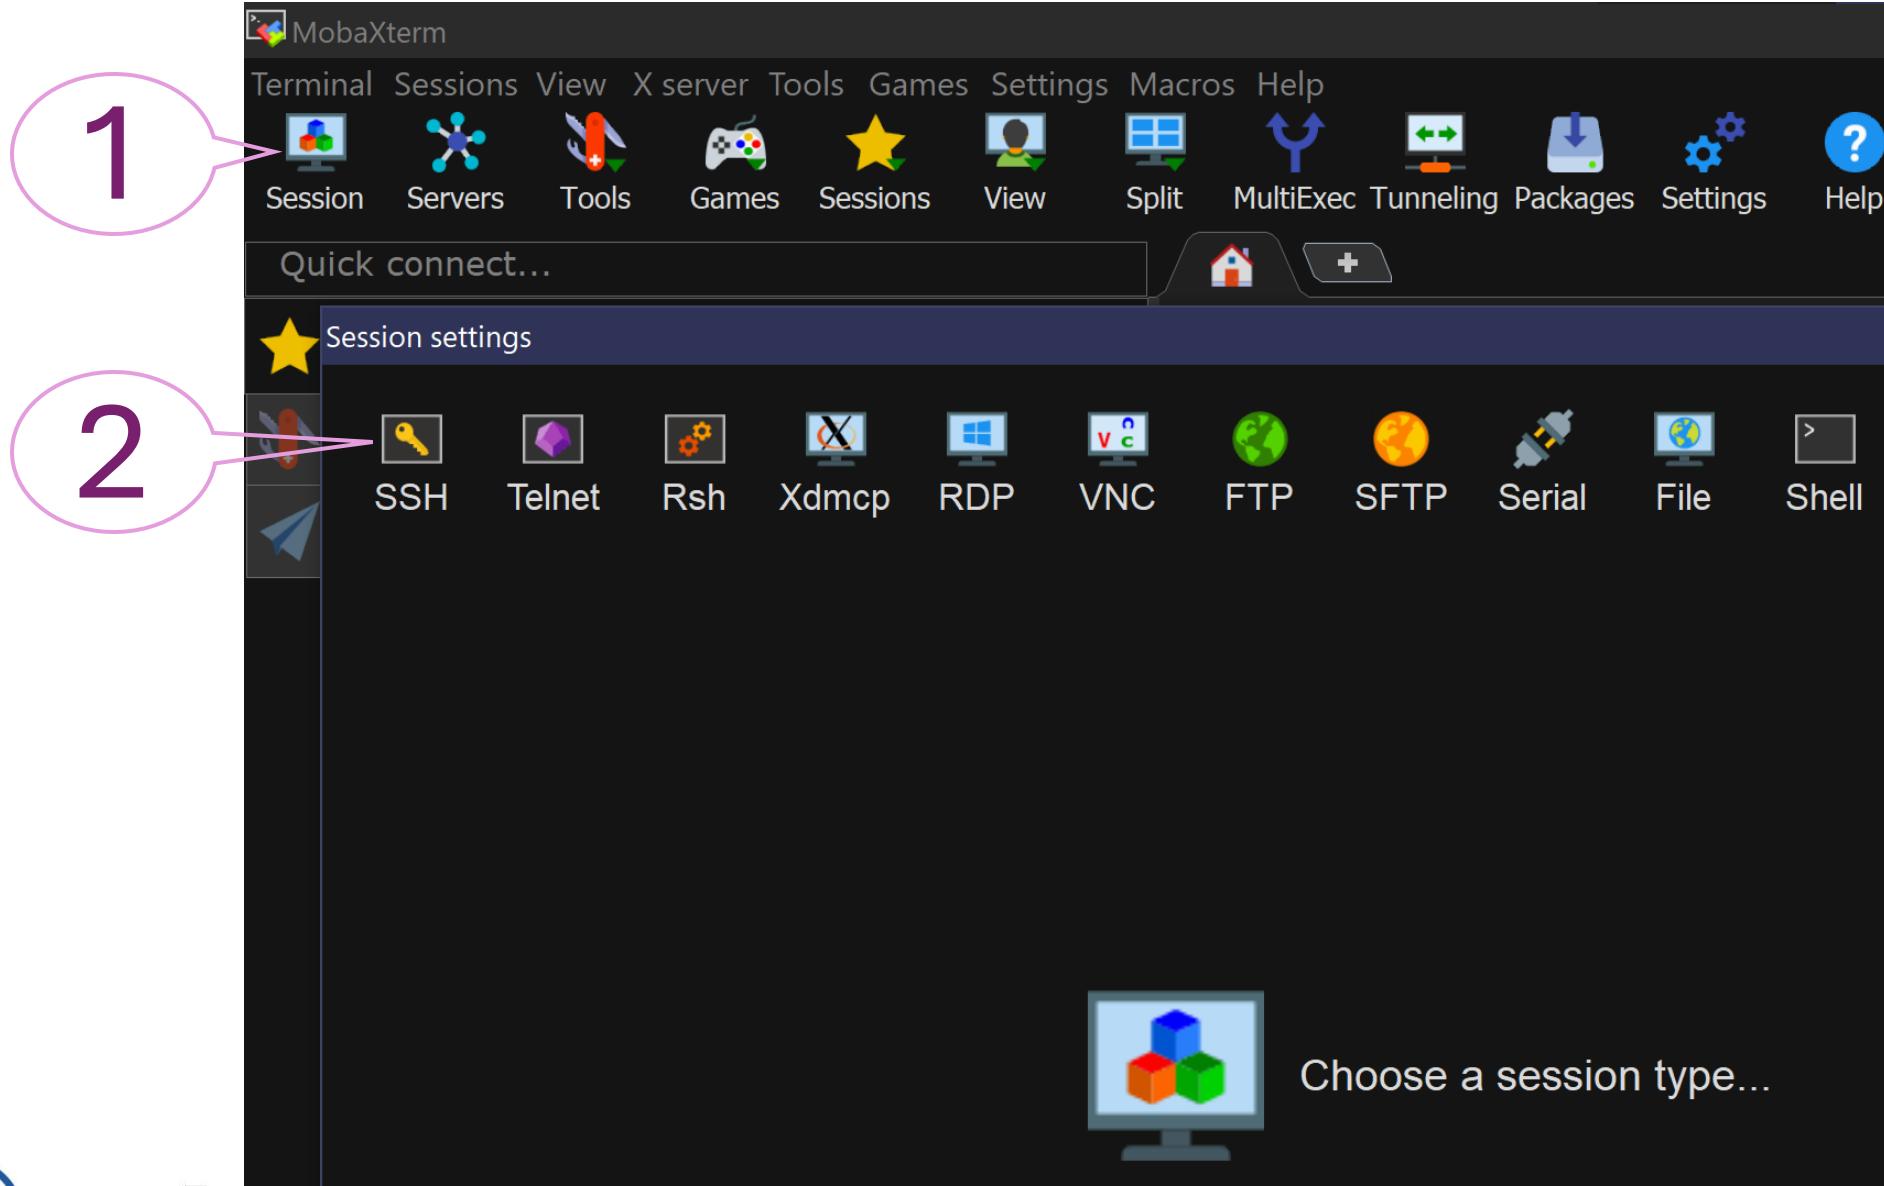
Margarida Madeira e Moura
Ilyass Hankrir

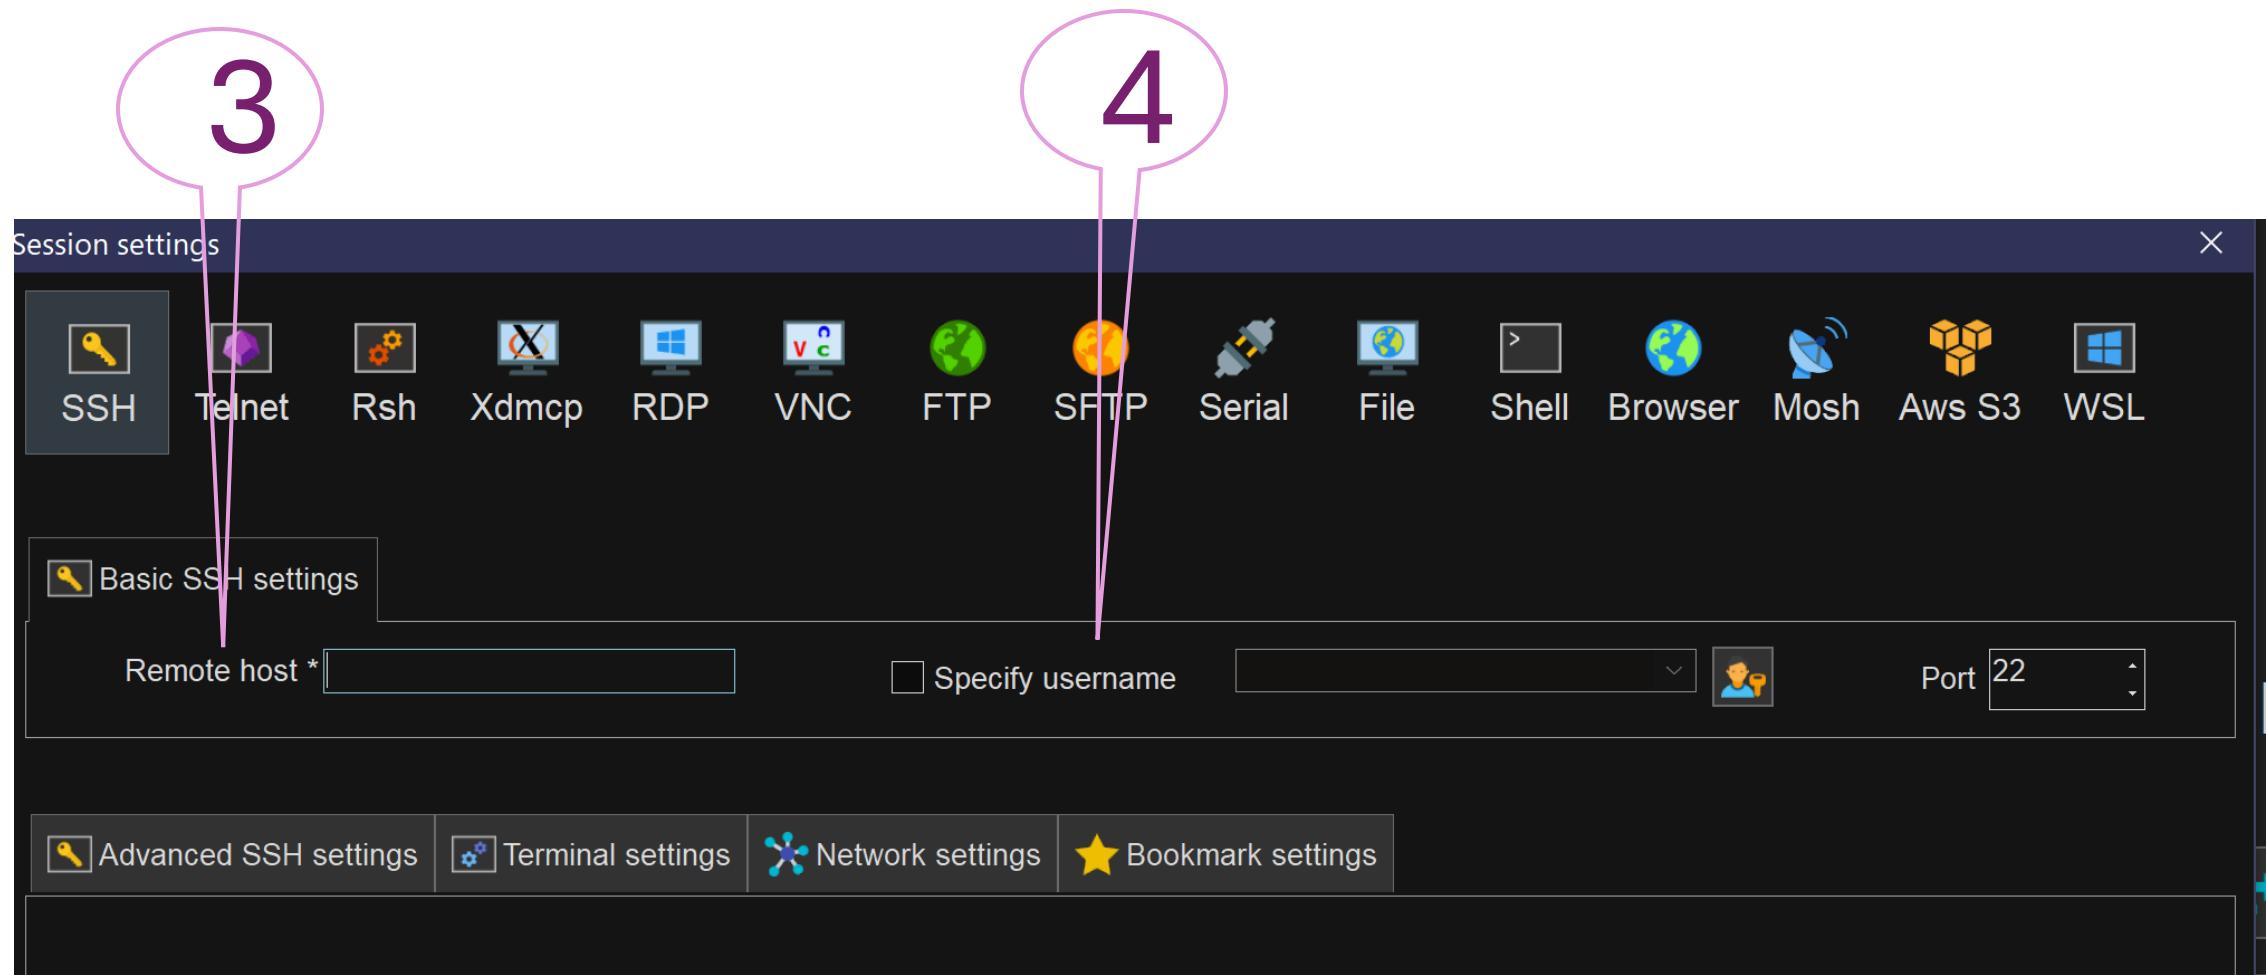
HPCvLAB@UALGARVE
April 11th, 2025

Tools

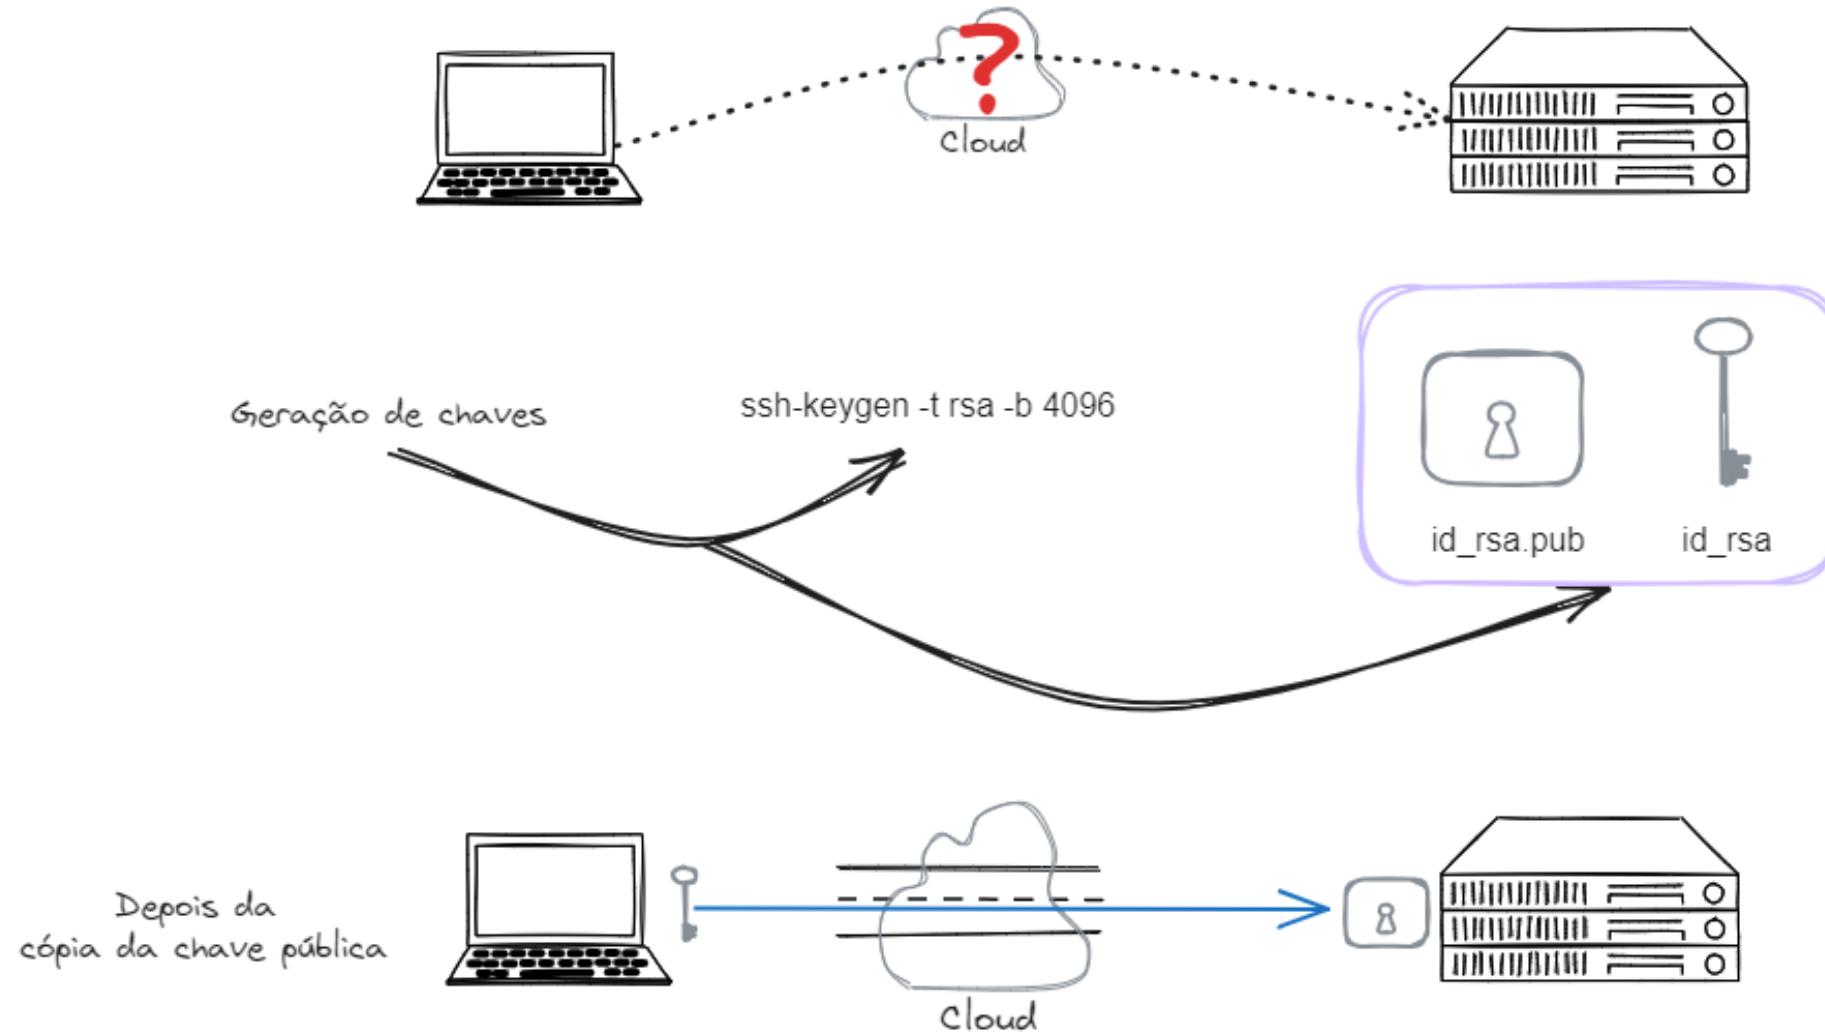
If using a Microsoft Windows OS:

- MobaXTerm Home Edition (available from <https://mobaxterm.mobatek.net/>)

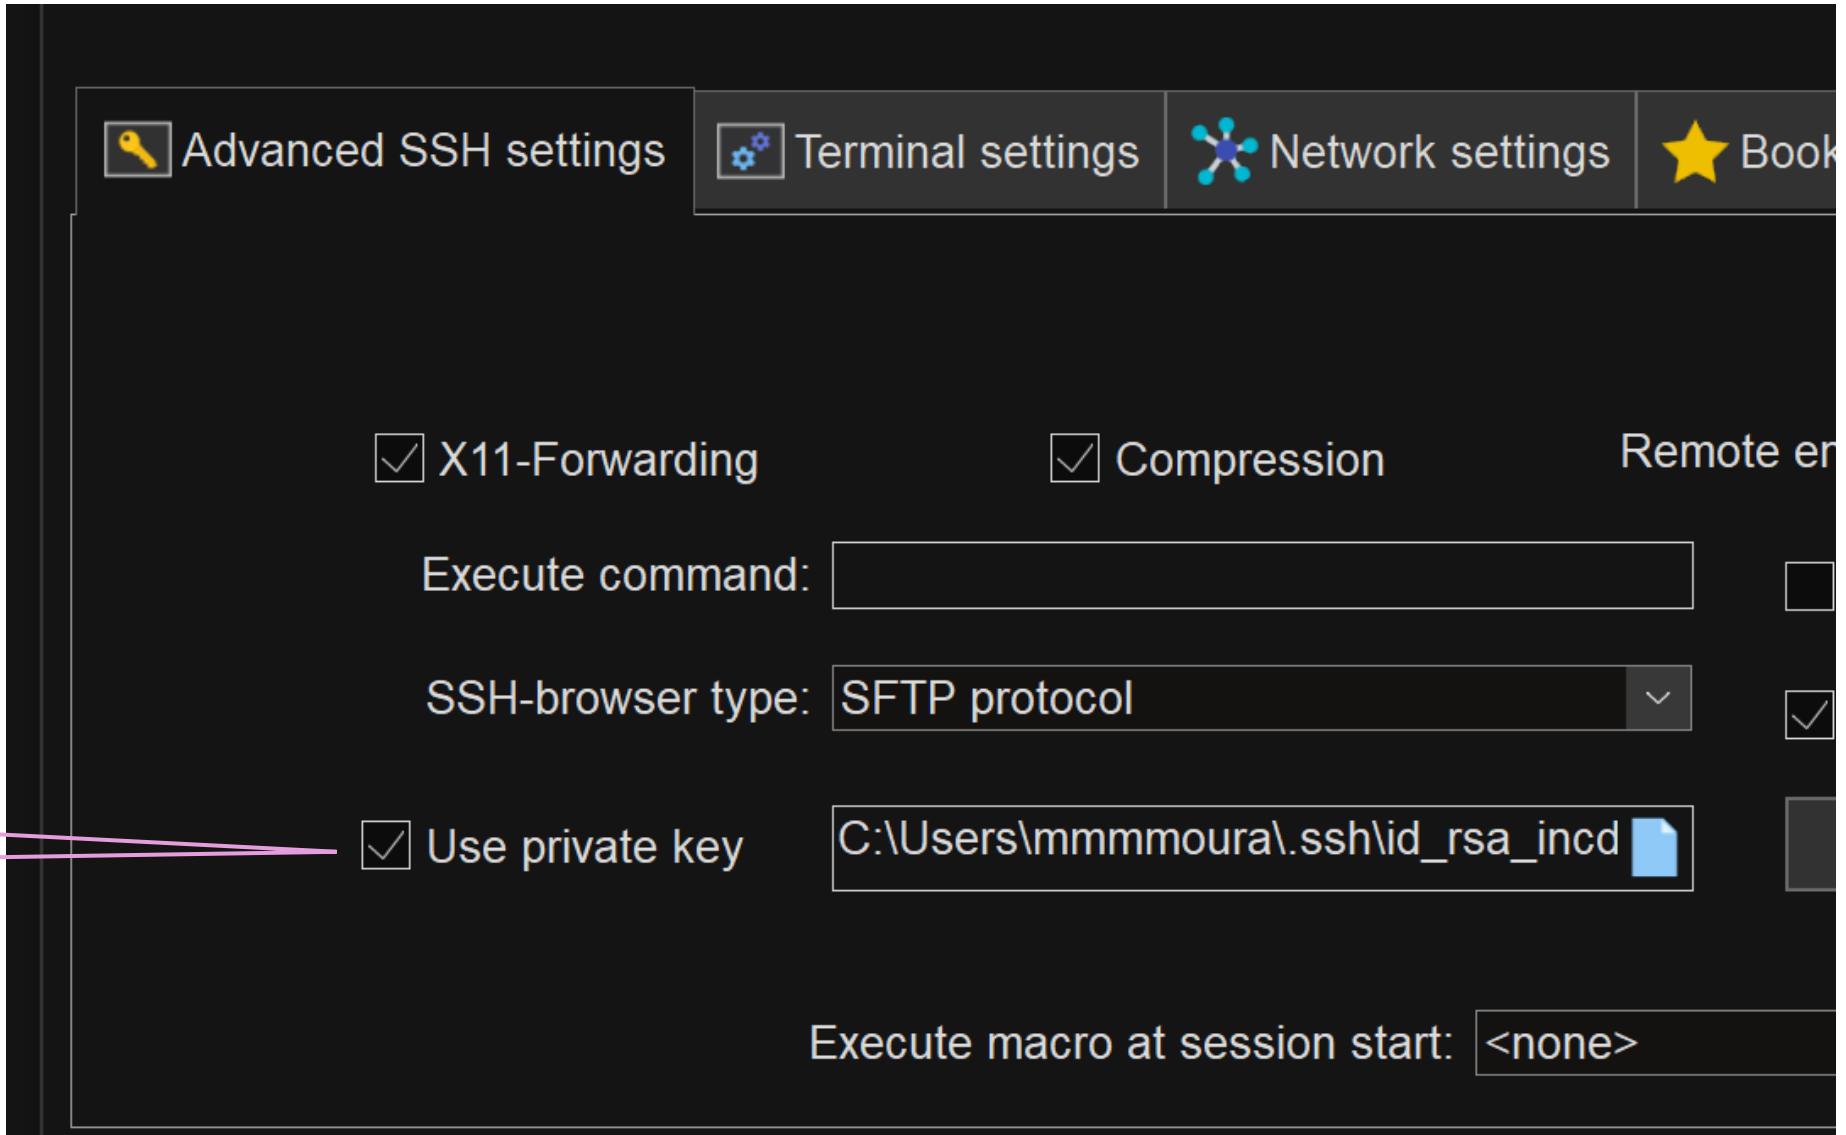




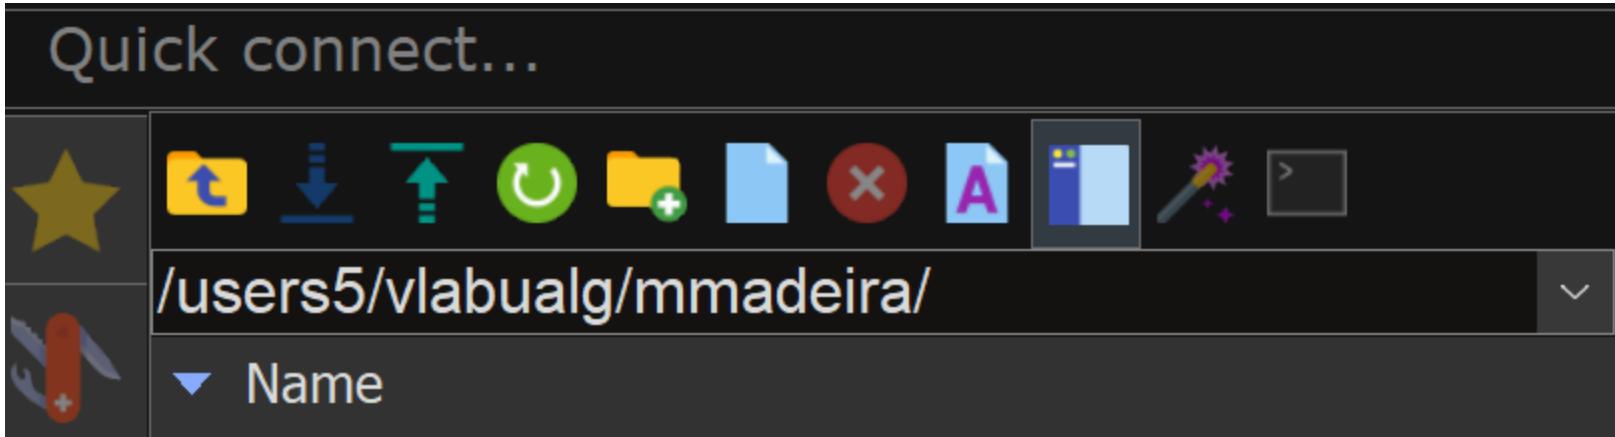
3 – cirrus.a.incd.pt 4 – the username



5



MobaXterm buttons



Tools

If using a Microsoft Windows OS:

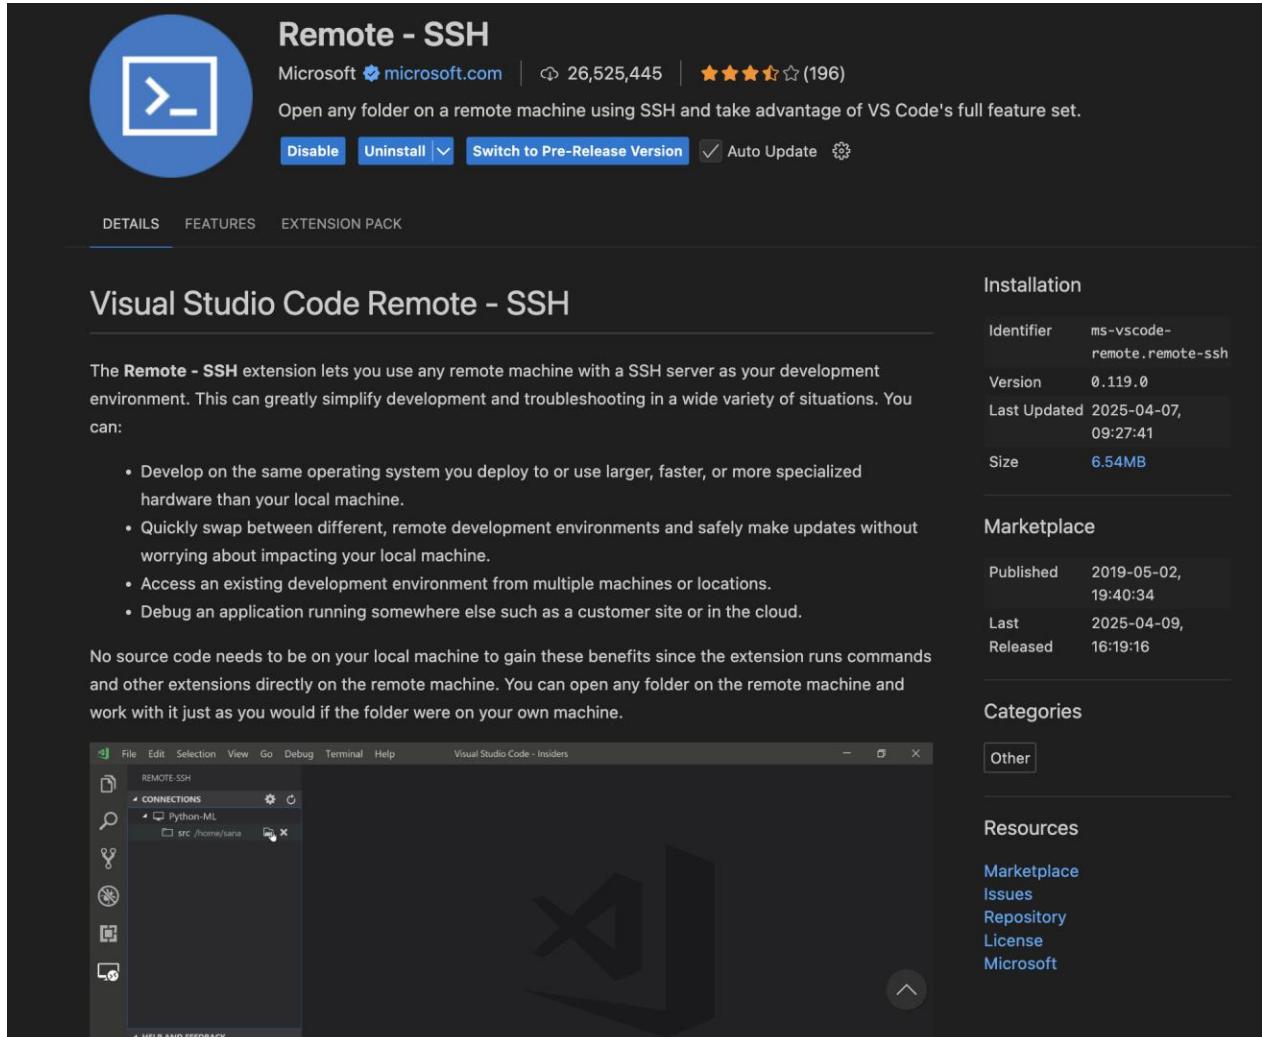
- MobaXTerm Home Edition (available from <https://mobaxterm.mobatek.net/>)

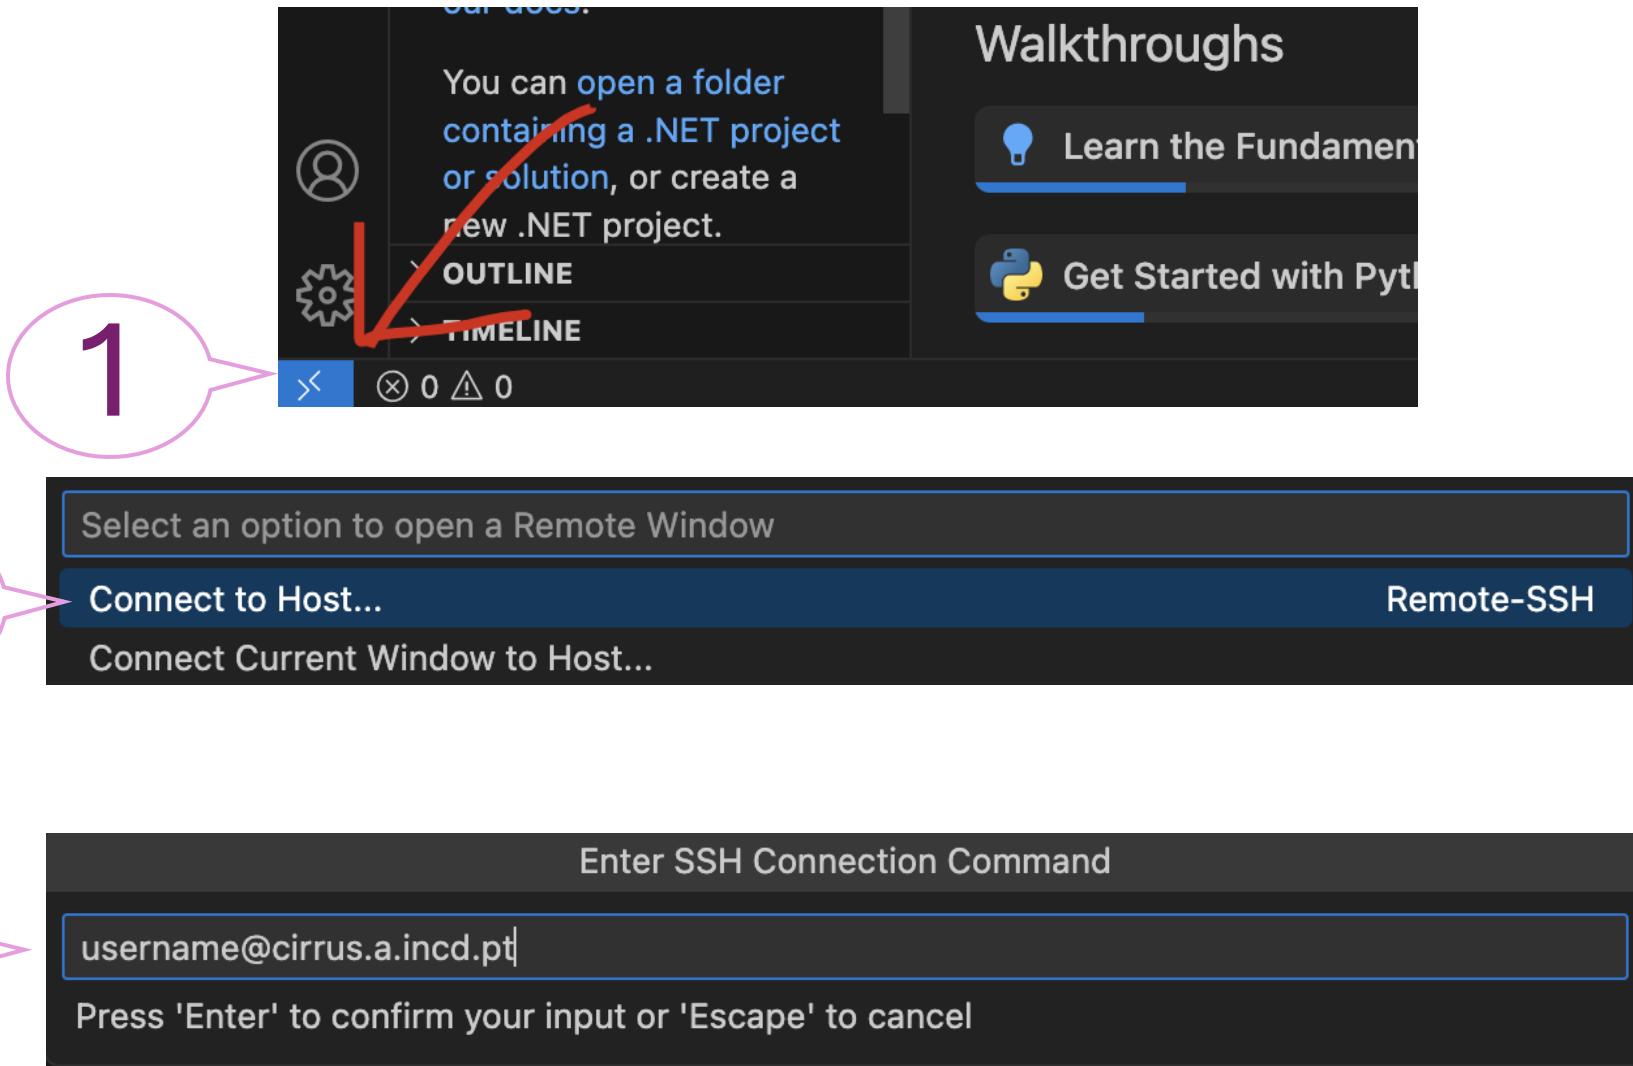
Alternatives for Windows OS/MacOS:

- Visual Studio Code

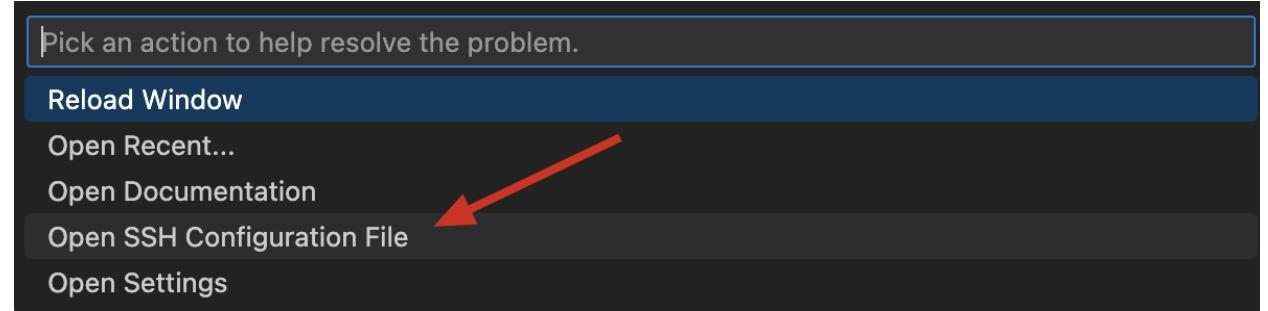
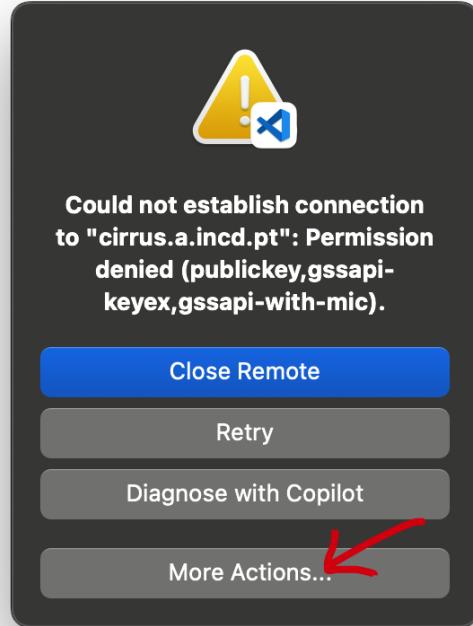
Starting a remote session using VS Code

- Install the following extension if not already installed



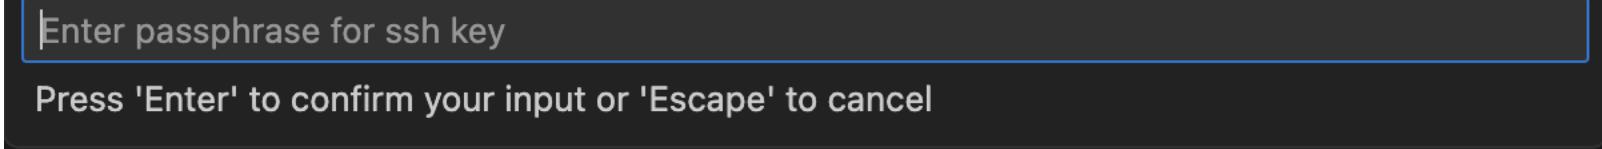


When encountering
this error follow
these steps:



Make sure your config includes:
IdentityFile ~/Path/To/Private_Key

```
1 Host cirrus.a.incd.pt
2 HostName cirrus.a.incd.pt
3 User username
4 IdentityFile ~/.ssh/id_rsa_incd
```



Tools

If using a Microsoft Windows OS:

- MobaXTerm Home Edition (available from <https://mobaxterm.mobatek.net/>)

Alternatives for Windows OS/MacOS:

- Visual Studio Code

If using the terminal:

Define the username if different
from your current user

```
~ — -zsh
user@PC ssh -i ~/.ssh/id_rsa_incd username@cirrus.a.incd.pt
Enter passphrase for key '/Users/.ssh/id_rsa_incd': 
```

After log in

```
#####
# NCG INGRID Public Login Service
# ingrid.helpdesk@lip.pt
#
# Auto logout after 12 hours with no activity
#
# https://wiki.incd.pt/books/manage-jobs/chapter/manage-slurm-jobs
#
#          Max
# PARTIT. NODES CORES NODETYPE
# hpc*      13 1728  AMD   EPYC 7643           (default )
# fct       13 1728  AMD   EPYC 7643           (reserved)
# short     18    36  AMD   EPYC 7643           (30 min. max.)
# gpu       2    192 AMD   EPYC 7643 + NVidia A100
# gpu       2    192 AMD   EPYC 7552 + NVidia Tesla V100s + Tesla T4
# gpu       2    192 AMD   EPYC 7552 + NVidia Tesla T4 + Tesla T4
#
# 01-04-2024: AlmaLinux 8 is now the official cluster
# old cluster user interface based on CentOS 7 available on:
# cirrus7.a.incd.pt
#
#####
```

Environment Management

Cirrus uses **environment modules** to manage compilers (GCC, Intel), MPI versions, Python, etc.

module avail**# show available modules**

module spider**# search for a module**

module load [module]**# load a module**

module list**# list all loaded modules**

module unload [module]**# unload a module**

module purge**# unload all modules**

Environment Management

module avail

```
----- /cvmfs/sw.el8/modules/hpc/intel -----
gmxMMPBSA/1.6.3          intel/libs/gsl/2.7           intel/libs/libpng/1.6.39      intel/netcdf-cxx4/4.3.1      intel/openmpi/4.1.4
intel/castep/24.1         intel/libs/hdf5/1.14.0        intel/libs/nlopt/2.7.1       intel/netcdf-fortran/4.6.1    intel/openmpi/4.1.6 (D)
intel/gamess/2023-R2       intel/libs/hdf5/1.14.3 (D)   intel/mvapich2/2.3.7        intel/nwchem/nwchem-7.2.2
intel/libsblas/3.11.0      intel/libs/jemalloc/5.3.0     intel/netcdf-c/4.9.2        intel/openfoam-org/11
intel/libfftw/3.3.10       intel/libs/lapack/3.11.0      intel/netcdf-cxx/4.2        intel/openfoam/2306

----- /cvmfs/sw.el8/modules/gpu -----
cuda/10.2      cuda/11.8      cuda/12.6      (D)      nvhpc-nompi/23.1      tensorflow/2.14.0
cuda/11.2      cuda/12.1      nvhpc-byo-compiler/23.1      nvhpc/23.1      tensorflow/2.18.0 (D)

----- /cvmfs/sw.el8/modules/bio -----
R/4.2.2          blast-plus/2.12.0      fastqc/0.12.1      (D)      htslib/1.17      (D)      picard/3.1.1      star/2.7.6a
R/4.4.3          blast-plus/2.14.1      figtree/0.12.1      (D)      igv/2.12.3      plink/1.07      star/2.7.10b (D)
admixture/1.3.0      bowtie2/2.4.2      figtree/1.4.3      (D)      iqtree2/2.1.2      plink2/2.00a4.3  stringtie/3.0.0
alphafold/2.3.2      bowtie2/2.5.1      finestructure/4.1.1    iqtree2/2.2.2      (D)      population2/1.201 suitesparse/7.8.2
alphapulldown/2.0.1      bwa/0.7.17      freebayes/1.3.6      jellyfish/2.3.1      raxml/8.2.12      transdecoder/5.5.0
autodock-gpu-develop/11.3.0      chimera/1.18      gatk/4.5.0.0      kallisto/0.50.1      salmon/1.10.3      transdecoder/5.7.1 (D)
autodock-vina/1.2.3      chromopainter/0.0.4      grenepipe/0.14.0      mrbayes/3.2.7a      samtools/1.8      trimgalore/0.6.10
bcftools/1.8          cutadapt/4.4      hisat2/2.2.1      openbabel/3.0.0      samtools/1.17      (D)      trinity/2.14.0
beagle/5.4          eigensoft/8.0.0      hmmer/3.3.2      orca/6.0.1      sickle/1.33      vcftools/0.1.14
bismark/0.23.0      evigene/23.7.15      hmmer/3.4      (D)      orca/shared/5.0.4  sratoolkit/3.0.0  vcftools/0.1.16 (D)
bismark/0.24.1      (D)          fastqc/0.11.9      htslib/1.8      orca/static/5.0.4  stAIcalc      vina-gpu/2.1

Where:
D: Default Module

If the avail list is too long consider trying:
"module --default avail" or "ml -d av" to just list the default modules.
"module overview" or "ml ov" to display the number of modules for each name.
```

Environment Management

module avail	# show available modules
module spider	# search for a module
module load [module]	# load a module
module list	# list all loaded modules
module unload [module]	# unload a module
module purge	# unload all modules

Environment Management

module spider

```
[username@cirrus $ module spider python
```

```
-----  
python:
```

```
-----  
Versions:
```

```
  python/3.7  
  python/3.8  
  python/3.10  
  python/3.11
```

```
-----  
For detailed information about a specific "python" package (including how to load the modules)  
use the module's full name.
```

```
Note that names that have a trailing (E) are extensions provided by other modules.
```

```
For example:
```

```
$ module spider python/3.11
```

Environment Management

module avail	# show available modules
module spider	# search for a module
module load [module]	# load a module
module list	# list all loaded modules
module unload [module]	# unload a module
module purge	# unload all modules

Environment Management

module load + module list

```
[username@cirrus $ module load python/3.10
[username@cirrus $ module list
```

Currently Loaded Modules:

1) gcc-11.3 2) gcc11/openmpi/4.1.4 3) python/3.10

Environment Management

module avail	# show available modules
module spider	# search for a module
module load [module]	# load a module
module list	# list all loaded modules
module unload [module]	# unload a module
module purge	# unload all modules

Environment Management

module unload

```
[username@cirrus $ module unload python/3.10  
[username@cirrus $ module list
```

Currently Loaded Modules:

- 1) gcc-11.3 2) gcc11/openmpi/4.1.4

Environment Management

module avail	# show available modules
module spider	# search for a module
module load [module]	# load a module
module list	# list all loaded modules
module unload [module]	# unload a module
module purge	# unload all modules

Environment Management

module purge

```
[username@cirrus $ module purge
[username@cirrus $ module list
No modules loaded
username@cirrus $ ]
```

Examples

Example 1 – distinguishing output from the job script and from the job load

Example 2 – what if something goes wrong: %j.err

Example 3 – MPI in C

Example 4 – OpenMP in C

Example 5 – MPI + OpenMP in C

Example 6 – GPU in C

Example 1 – Hello HPC World

The image shows two code editors side-by-side. The left editor, titled 'hello.sh', contains a bash script. The right editor, titled 'hello.py', contains a Python script.

```
hello.sh content:
1  #!/bin/bash
2
3  ## If you want to be updated by mail
4  #SBATCH --mail-user=your_user_name@your_mailserver
5  #SBATCH --mail-type=ALL
6
7  #SBATCH --job-name=hello
8  #SBATCH --partition=hpc
9  #SBATCH --qos=cpuvlabualg
10 #SBATCH --ntasks=1
11
12 ## The result of the job,
13 ## |respectively stdout e stderr
14 #SBATCH --output=%x.%j.out
15 #SBATCH --error=%x.%j.err
16
17 module load python
18 echo "Hello HPC World, from bash script!"
19 python3 hello.py
```

```
hello.py content:
1  #!/usr/bin/python3
2  # -*- coding: utf-8 -*-
3
4  """
5      (c) Margarida Madeira e Moura, 2024
6      Demo just_hello
7
8  """
9  print("Hello HPC World, from Python!")
```

Submit your job:
\$ sbatch hello.sh

Example 1 – Hello HPC World.out

```
[username@cirrus $ cat hello.21772859.out
* -----
* Running PROLOG for hello on Tue May 6 15:07:06 WEST 2025
*   JOB_NAME          : hello
*   JOB_ID           : 21772859
*   JOB_PARTITION    : hpc
*   JOB_USER          : username
*   JOB_ACCOUNT       : vlabualg
*   JOB_QOS           : normal
*   NODE_LIST          : hpc071
*   SLURM_NNODES      : 1
*   SLURM_NPROCS       : 1
*   SLURM_NTASKS       : 1
*   SLURM_CPUS_ON_NODE : 1
*   SLURM_TASKS_PER_NODE : 1
*   SLURM_JOB_CPUS_PER_NODE : 1
*   SLURM_MEM_PER_CPU  : 5000
*   SUBMIT_HOST        : cirrus.a.incd.pt
*   WORK_DIR           : /users5/vlabualg/examples/just_hello
*   JOB_SCRIPT         : /users5/vlabualg/examples/just_hello/hello.sh
* -----
Hello HPC World, from bash script!
Hello HPC World, from Python!
username@cirrus $ ]
```

Example 2 – Hello HPC World.err

```
[username@cirrus $ cat hello.21772948.out
* -----
* Running PROLOG for hello on Tue May 13 05:47:13 WEST 2025
*   JOB_NAME          : hello
*   JOB_ID            : 21772948
*   JOB_PARTITION     : hpc
*   JOB_USER          : username
*   JOB_ACCOUNT       : vlabualg
*   JOB_QOS           : normal
*   NODE_LIST          : hpc071
*   SLURM_NNODES      : 1
*   SLURM_NPROCS       : 1
*   SLURM_NTASKS       : 1
*   SLURM_CPUS_ON_NODE: 1
*   SLURM_TASKS_PER_NODE: 1
*   SLURM_JOB_CPUS_PER_NODE: 1
*   SLURM_MEM_PER_CPU  : 5000
*   SUBMIT_HOST        : cirrus09.a.incd.pt
*   WORK_DIR           : /users5/vlabualg/username/examples/just_hello
*   JOB_SCRIPT         : /users5/vlabualg/username/examples/just_hello/hello.sh
* -----
Hello HPC World, from bash script!
username@cirrus $ ]
```

hello.id.out

```
[username@cirrus $ cat hello.21772948.err
      File "/users5/vlabualg/examples/just_hello/hello.py", line 9
          print("Hello HPC World, from Python!")
                           ^
SyntaxError: unterminated string literal (detected at line 9)
username@cirrus $ ]
```

hello.id.err

sacct

[username@cirrus \$ sacct							
JobID	JobName	Partition	Account	AllocCPUS	State	ExitCode	
21772859	hello	hpc	vlabualg	1	COMPLETED	0:0	
21772859.ba+	batch		vlabualg	1	COMPLETED	0:0	
21772859.ex+	extern		vlabualg	1	COMPLETED	0:0	

sacct

username@cirrus08 just_hello]\$ sacct --format=User,JobID,Jobname,partition,state,time,start,end,elapsed,MaxRSS,MaxVMSize,nnodes,ncpus,nodelist -u username --start 2024-12-10													
User	JobID	JobName	Partition	State	Timelimit	Start	End	Elapsed	MaxRSS	MaxVMSize	NNodes	NCpus	NodeList
username	19093687	hello	hpc	COMPLETED	4:00:00:00	2024-12-10T16:13:15	2024-12-10T16:13:16	00:00:01			1	1	hpc067
	19093687.ba+	batch		COMPLETED		2024-12-10T16:13:15	2024-12-10T16:13:16	00:00:01	0	956K	1	1	hpc067
	19093687.ex+	extern		COMPLETED		2024-12-10T16:13:15	2024-12-10T16:13:16	00:00:01	0	0	1	1	hpc067
	19093687.0	python3		COMPLETED		2024-12-10T16:13:16	2024-12-10T16:13:16	00:00:00	0	4K	1	1	hpc067
username	19093704	hello	hpc	COMPLETED	4:00:00:00	2024-12-10T16:22:57	2024-12-10T16:22:58	00:00:01			1	1	hpc067
	19093704.ba+	batch		COMPLETED		2024-12-10T16:22:57	2024-12-10T16:22:58	00:00:01	0	4K	1	1	hpc067
	19093704.ex+	extern		COMPLETED		2024-12-10T16:22:57	2024-12-10T16:22:58	00:00:01	0	0	1	1	hpc067
username	19093715	brute_for+	hpc	COMPLETED	4:00:00:00	2024-12-10T16:31:07	2024-12-10T16:31:10	00:00:03			1	1	hpc067
	19093715.ba+	batch		COMPLETED		2024-12-10T16:31:07	2024-12-10T16:31:10	00:00:03	0	20K	1	1	hpc067
	19093715.ex+	extern		COMPLETED		2024-12-10T16:31:07	2024-12-10T16:31:10	00:00:03	0	0	1	1	hpc067

To list the waiting queue, use:
squeue

To cancel a job, use:
scancel <JOBID>

sacct

username@cirrus08 just_hello]\$ sacct --format=User,JobID,Jobname,partition,state,time,start,end,elapsed,MaxRSS,MaxVMSize,nnodes,ncpus,nodelist -u username --start 2024-12-10													
User	JobID	JobName	Partition	State	Timelimit	Start	End	Elapsed	MaxRSS	MaxVMSize	NNodes	NCPUS	NodeList
username	19093687	hello	hpc	COMPLETED	4-00:00:00	2024-12-10T16:13:15	2024-12-10T16:13:16	00:00:01			1	1	hpc067
	19093687.ba+	batch		COMPLETED		2024-12-10T16:13:15	2024-12-10T16:13:16	00:00:01	0	956K	1	1	hpc067
	19093687.ex+	extern		COMPLETED		2024-12-10T16:13:15	2024-12-10T16:13:16	00:00:01	0	0	1	1	hpc067
	19093687.0	python3		COMPLETED		2024-12-10T16:13:16	2024-12-10T16:13:16	00:00:00	0	4K	1	1	hpc067
username	19093704	hello	hpc	COMPLETED	4-00:00:00	2024-12-10T16:22:57	2024-12-10T16:22:58	00:00:01			1	1	hpc067
	19093704.ba+	batch		COMPLETED		2024-12-10T16:22:57	2024-12-10T16:22:58	00:00:01	0	4K	1	1	hpc067
	19093704.ex+	extern		COMPLETED		2024-12-10T16:22:57	2024-12-10T16:22:58	00:00:01	0	0	1	1	hpc067
username	19093715	brute_for+	hpc	COMPLETED	4-00:00:00	2024-12-10T16:31:07	2024-12-10T16:31:10	00:00:03			1	1	hpc067
	19093715.ba+	batch		COMPLETED		2024-12-10T16:31:07	2024-12-10T16:31:10	00:00:03	0	20K	1	1	hpc067
	19093715.ex+	extern		COMPLETED		2024-12-10T16:31:07	2024-12-10T16:31:10	00:00:03	0	0	1	1	hpc067

To list the waiting queue, use:
squeue

username@cirrus \$ squeue							
JOBID	PARTITION	NAME	USER	ST	TIME	NODES	CPUS
21772922	hpc	hello	username	PD	0:00	1	4

To cancel a job, use:
scancel <JOBID>

Parallel Programming Models: MPI vs OpenMP

Feature	OpenMP	MPI
Memory Model	Shared memory (same node)	Distributed memory (multi-node)
Typical Use	Multithreading	Multiprocessing
Syntax	Pragmas / directives	Function calls
Communication	Implicit (via memory)	Explicit (via messages)
Cirrus Usage Example	--cpus-per-task=4	--ntasks=4

Example 3 – Hello_MPI

```
#include <mpi.h>
#include <stdio.h>

int main(int argc, char** argv) {
    MPI_Init(NULL, NULL);

    int world_size;
    MPI_Comm_size(MPI_COMM_WORLD, &world_size);

    int world_rank;
    MPI_Comm_rank(MPI_COMM_WORLD, &world_rank);

    printf("Hello from process %d of %d\n", world_rank, world_size);

    MPI_Finalize();
}
```

```
1  #!/bin/bash
2  ## If you want to be updated by mail
3  #SBATCH --mail-user=youremail@example.com
4  #SBATCH --mail-type=ALL
5
6  #SBATCH --job-name=hello
7  #SBATCH --partition=hpc      # Partition to submit to
8  #SBATCH --nodes=1            # Number of nodes
9  #SBATCH --ntasks=4          # Total number of MPI tasks
10 #SBATCH --output=%x.%j.out
11 #SBATCH --error=%x.%j.err
12
13 module load gcc11/openmpi/4.1.4
14
15 mpicc -o hello_mpi hello_mpi.c
16
17 mpirun ./hello_mpi
```

Submit your job:

\$ sbatch hello_mpi.sh

Example 3 – Hello_MPI

```
* -----
* Running PROLOG for hello on Thu Apr 10 07:48:02 WEST 2025
*   JOB_NAME          : hello
*   JOB_ID            : 21003264
*   JOB_PARTITION     : hpc
*   JOB_USER          : username
*   JOB_ACCOUNT       : vlabualg
*   JOB_QOS           : normal
*   NODE_LIST          : hpc080
*   SLURM_NNODES      : 1
*   SLURM_NPROCS       : 4
*   SLURM_NTASKS       : 4
*   SLURM_CPUS_ON_NODE : 4
*   SLURM_TASKS_PER_NODE : 4
*   SLURM_JOB_CPUS_PER_NODE : 4
*   SLURM_MEM_PER_CPU    : 5000
*   SUBMIT_HOST         : cirrus.a.incd.pt
*   WORK_DIR           : /users5/vlabualg/username/examples
*   JOB_SCRIPT          : /users5/vlabualg/username/examples/hello_mpi.sh
*
Hello from process 3 of 4
Hello from process 0 of 4
Hello from process 1 of 4
Hello from process 2 of 4
```

Example 4 – Hello_OpenMP

```
#!/bin/bash
## If you want to be updated by mail
#SBATCH --mail-user=youremail@example.com
#SBATCH --mail-type=ALL

#SBATCH --job-name=omp_hello
#SBATCH --partition=hpc
#SBATCH --qos=normal
#SBATCH --nodes=1
#SBATCH --ntasks=1
#SBATCH --cpus-per-task=4
#SBATCH --output=%x.%j.out
#SBATCH --error=%x.%j.err

module load gcc-11.3          # Load GCC with OpenMP support

export OMP_NUM_THREADS=$SLURM_CPUS_PER_TASK

gcc -fopenmp -o omp_hello omp_hello.c

# Run the program
./omp_hello
```

Submit your job:
\$ sbatch omp_hello.sh

```
#include <omp.h>
#include <stdio.h>
#include <stdlib.h>

int main (int argc, char *argv[])
{
    int nthreads, tid;

    /* Fork a team of threads giving them their own copies of variables */
#pragma omp parallel private(nthreads, tid)
    {
        /* Obtain thread number */
        tid = omp_get_thread_num();
        printf("Hello World from thread = %d\n", tid);

        /* Only master thread does this */
        if (tid == 0)
        {
            nthreads = omp_get_num_threads();
            printf("Number of threads = %d\n", nthreads);
        }

    } /* All threads join master thread and disband */
}
```

Example 4 – Hello_OpenMP

```
* -----
* Running PROLOG for omp_hello on Fri Apr 11 05:02:46 WEST 2025
*   JOB_NAME          : omp_hello
*   JOB_ID            : 21013153
*   JOB_PARTITION     : hpc
*   JOB_USER          : username
*   JOB_ACCOUNT       : vlabualg
*   JOB_QOS           : normal
*   NODE_LIST          : hpc070
*   SLURM_NNODES      : 1
*   SLURM_NPROCS       : 1
*   SLURM_NTASKS       : 1
*   SLURM_CPUS_ON_NODE : 4
*   SLURM_TASKS_PER_NODE : 1
*   SLURM_JOB_CPUS_PER_NODE : 4
*   SLURM_MEM_PER_CPU    : 5000
*   SUBMIT_HOST         : cirrus08.a.incd.pt
*   WORK_DIR           : /users/vlabualg/username/examples/OpenMP
*   JOB_SCRIPT          : /users/vlabualg/username/examples/OpenMP/omp_hello.sh
* -----
Hello World from thread = 0
Number of threads = 4
Hello World from thread = 2
Hello World from thread = 3
Hello World from thread = 1
```

Example 5 – Hello_Hybrid

```
#include <mpi.h>
#include <omp.h>
#include <stdio.h>

int main(int argc, char *argv[]) {
    int rank, size;

    // Initialize MPI
    MPI_Init(&argc, &argv);
    MPI_Comm_rank(MPI_COMM_WORLD, &rank); // process ID
    MPI_Comm_size(MPI_COMM_WORLD, &size); // total number of processes

    // Parallel OpenMP region inside each MPI process
    #pragma omp parallel
    {
        int tid = omp_get_thread_num();                  // thread ID
        int nthreads = omp_get_num_threads();           // threads in this process

        printf("Hello from thread %d out of %d in MPI process %d out of %d\n",
               tid, nthreads, rank, size);
    }

    MPI_Finalize(); // Finalize MPI
    return 0;
}
```

Example 5 – Hello_Hybrid

```
#!/bin/bash
## If you want email notifications
#SBATCH --mail-user=youremail@example.com
#SBATCH --mail-type=ALL

#SBATCH --job-name=hybrid_hello          # Job name
#SBATCH --partition=hpc                 # Partition
#SBATCH --qos=normal                   # Adjust to valid QoS for your project
#SBATCH --nodes=1                       # Number of nodes
#SBATCH --ntasks=2                      # Number of MPI processes
#SBATCH --cpus-per-task=4               # Threads per MPI process (OpenMP)
#SBATCH --output=%x.%j.out              # Standard output
#SBATCH --error=%x.%j.err               # Standard error

module load gcc11/openmpi/4.1.4          # Load MPI + OpenMP-capable compiler

export OMP_NUM_THREADS=$SLURM_CPUS_PER_TASK

echo "Compiling hybrid hello with mpicc and OpenMP support..."
mpicc -fopenmp -o hello_hybrid hello_hybrid.c

echo "Running hybrid hello..."
mpirun ./hello_hybrid
```

Submit your job:
\$ sbatch hello_hybrid.sh

Example 5 – Hello_Hybrid

```
* -----
* Running PROLOG for hybrid_hello on Fri Apr 11 05:20:30 WEST 2025
*   JOB_NAME          : hybrid_hello
*   JOB_ID           : 21013306
*   JOB_PARTITION    : hpc
*   JOB_USER          : username
*   JOB_ACCOUNT       : vlabualg
*   JOB_QOS           : normal
*   NODE_LIST          : hpc070
*   SLURM_NNODES      : 1
*   SLURM_NPROCS       : 2
*   SLURM_NTASKS       : 2
*   SLURM_CPUS_ON_NODE : 8
*   SLURM_TASKS_PER_NODE : 2
*   SLURM_JOB_CPUS_PER_NODE : 8
*   SLURM_MEM_PER_CPU  : 5000
*   SUBMIT_HOST        : cirrus08.a.incd.pt
*   WORK_DIR           : /users/vlabualg/username/examples/hybrid
*   JOB_SCRIPT         : /users/vlabualg/username/examples/hybrid/hello_hybrid.sh
* -----
Compiling hybrid hello with mpicc and OpenMP support...
Running hybrid hello...
Hello from thread 0 out of 4 in MPI process 0 out of 2
Hello from thread 2 out of 4 in MPI process 0 out of 2
Hello from thread 0 out of 4 in MPI process 1 out of 2
Hello from thread 2 out of 4 in MPI process 1 out of 2
Hello from thread 3 out of 4 in MPI process 1 out of 2
Hello from thread 1 out of 4 in MPI process 1 out of 2
Hello from thread 3 out of 4 in MPI process 0 out of 2
Hello from thread 1 out of 4 in MPI process 0 out of 2
```

Example 6 – Hello_GPU

```
#include <stdio.h>

__device__ const char *STR = "HELLO WORLD!";
const char STR_LENGTH = 12;

__global__ void hello()
{
    printf("%c\n", STR[threadIdx.x % STR_LENGTH]);
}

int main(void)
{
    int num_threads = STR_LENGTH;
    int num_blocks = 1;
    hello<<<num_blocks,num_threads>>>();
    cudaDeviceSynchronize();
    return 0;
}
```

```
#!/bin/bash

#SBATCH --mail-user=youremail@example.com
#SBATCH --mail-type=ALL

#SBATCH --job-name=hello00
#SBATCH --partition=gpu
#SBATCH --qos=gpuvlabualg

#SBATCH --gres=gpu

#SBATCH --output=%x.%j.out
#SBATCH --error=%x.%j.err

module purge
module load cuda/12.6
nvcc -o hello hello.cu

echo `hostname` 
srun ./hello
```

Submit your job:
\$ sbatch hello.sh

Example 6 – Hello_GPU

```
* ---  
* Running PROLOG for hello00 on Mon Nov 25 19:41:43 WET 2024  
*   JOB_NAME          : hello00  
*   JOB_ID           : 18969544  
*   JOB_PARTITION     : gpu  
*   JOB_USER          : ilyass  
*   JOB_ACCOUNT       : vlabualg  
*   JOB_QOS           : gpubvlabualg  
*   JOB_GPUS          : 0  
*   NODE_LIST         : hpc060  
*   SLURM_NNODES      : 1  
*   SLURM_CPUS_ON_NODE: 1  
*   SLURM_TASKS_PER_NODE: 1  
*   SLURM_JOB_CPUS_PER_NODE: 1  
*   SLURM_GPUS_ON_NODE: 1  
*   SUBMIT_HOST        : cirrus09.a.incd.pt  
*   WORK_DIR          : /users5/vlabualg/ilyass/GPU  
*   JOB_SCRIPT         : /users5/vlabualg/ilyass/GPU/hello.sh  
*  
* GPU 0              : Name          Tesla V100S-PCIE-32GB  
*                      : SN            1421220029960  
*                      : Driver         560.35.03  
*                      : Memory         32494 MBytes  
*                      : Multiprocessors 80  
*                      : Cores/Multiprocessor 64  
*                      : Cores          5120  
* ---
```

```
hpc060.a.incd.pt  
H  
E  
L  
L  
0  
W  
0  
R  
L  
D  
!  
!
```

References

Software used in the presentation:

- mobaxterm, vscode
- Environment software
- Python
- Mpi
- Openmp
- CUDA

Codes from:

- <https://developer.nvidia.com/cuda-education>
- <https://hpc-tutorials.llnl.gov/>

Acknowledgements

This initiative is supported by the project “NATIONAL COMPETENCE CENTRES IN THE FRAMEWORK OF EUROHPC PHASE 2 – EUROCC2”, funded by Fundação para a Ciência e Tecnologia, I.P., and “DIGITAL-EUROHPC-JU-2022-NCC-01” da European High-Performance Computing Joint Undertaking ('JU') and makes use of Cirrus supercomputer of CNCA - National Distributed Computing Infrastructure.

Examples adapted from LLNL HPC Tutorials