



The THINC Cluster

Robert Klein*

December 4, 2023

Introduction

The current THINC cluster uses the SLURM system to schedule and run jobs. It currently runs Debian 11 and you have access to your “POLY” home directory and the “bee8” data storage.

The THINC cluster documentation is at <https://max.mpg.de/sites/poly/Research/Experts/Pages/HPC-Cluster.aspx>.

Please also take a look at the official SLURM documentation at <https://slurm.schedmd.com/documentation.html>.

A “Quick Reference Card” is at <https://slurm.schedmd.com/pdfs/summary.pdf>.

Preparation before using the cluster

Please set up your SSH keys as detailed in the welcome sheet, so you can copy your data using `scp` and/or `rsync`.

Keep in mind

- Write the result data to `/usr/scratch`. This directory is not backed up. Please remember to copy the result at the end of your job away (e.g. using the `rsync` command).
- Get information about available queues and their respective time limit using the `sinfo` command.

Best Practices

- jobs should write to `/usr/scratch/$LOGNAME`.
- copy the output to your personal disk space (e.g. on bee14) after the job has finished.

*Robert.Klein@mpip-mainz.mpg.de

Submit Script examples

Simple SLURM script

```
#!/bin/sh
printf "Hello_world.\n"
```

more complex SLURM script

```
#!/bin/bash -l

#          ( )
# SLURM Options < ) >
##### / \

# Define the partition on which the job will run. Defaults to
# CPU_Std32 if omitted.
# Partitions currently (December 2023) are:
# - CPU_Std20
# - CPU_Std32
# - CPU_IBm32
# - GPU_Std16
#SBATCH --partition=CPU_Std20

# Define, how many nodes you need. Here, we ask for 1 node.
# Only the COU_IBm32 partition ca use mor then one node
#SBATCH --nodes=1

# Number of cores (i.e. 'rank' in MPI) (defaults to 1, if omitted):
#SBATCH --ntasks=20

# mails? and to whom
#SBATCH --mail-type=END,FAIL
# SBATCH --mail-user=YOUR_USERNAME@mpip-mainz.mpg.de

#####
# no bash commands above this line
# all sbatch directives need to be placed before the first bash command
```

```
mpirun -np 20 ./a.out
```

some SLURM pointers

- The shell in the first line of your job script should be the same as your login shell. If you use any other shell, your job will be limited to the interactive time limit (15 min).
- The system's openmpi is not compiled with SLURM support, so you can't start jobs using "srun". (use "mpirun" instead)
- If you're using/initializing the intel compiler, your job script must be a bash script. First line should be:

```
#!/bin/bash
```

Building Software with cmake

- parallelization for CMAKE is set with `CMAKE_BUILD_PARALLEL_LEVEL`. You can set this to `$SLURM_NTASKS` in the SLURM script (`$SLURM_NTASKS` is the number of cores assigned to the SLURM job).