Max-Planck-Institut für Polymerforschung

Max Planck Institute for Polymer Research



The THINC Cluster

Robert Klein*

February 24, 2025

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.

Important Slurm commands

sbatch Submit a batch script to Slurm. *squeue* View information about jobs located in the Slurm scheduling queue. *scancel* Used to signal jobs or job steps that are under the control of Slurm. *sinfo* View information about Sluem nodes and partitions.

Submit Script examples

Simple Slurm script

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

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

more complex Slurm script

```
#!/bin/bash -1
#
                 (●_●)
# Slurm Options
                 <))>
################################
# Define the partition on which the job will run. Defaults to
# CPU_Std20 if omitted.
# Partitions currently (February 2025) are:
# - CPU_Std20
# - CPU_Std32
# - CPU_IBm32
# - GPU_Std16
#SBATCH --partition=CPU_Std32
# 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 32 ./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).