

Exercises for SLURM scripting

1) Log into tombo and obtain cluster details

5min: obtain the cluster information from slurm (command in presentation)

- how many nodes are available, in use, etc.

5min: obtain the state of the queue

```
$ sinfo
$ squeue
```

2) Obtain the environment variable set by SLURM in your job scope, then try to find their meaning using “man srun” and “man sbatch”. Note that those variable are available to your computation.

10min ~ 15min: (maybe difficult)

```
$ env > before.env
$ srun env > after.env
$ diff before.env after.env
```

3) Get verbose information about a running job job:

5 min: command was in the presentation

```
$ vi job_ex.slurm
#!/bin/bash
#SBATCH -J job_dep
#SBATCH -p compute
#SBATCH -n 1
#SBATCH -mem=100m
echo "${0} is running ..."
sleep 80
echo "${0} is still running ..."
sleep 60
echo "${0} is almost done ..."
sleep 50
echo "${0} will end"
$ sbatch job1.slurm
Submitted batch job <job_number>
$ scontrol show job <job_number>
```

4) Use job dependencies to run job2.slurm after job1.slurm have exited successfully

10 min: use man sbatch to find about “job” “dependency”

5 min: write script and run it

```
$ vi job1.slurm
#!/bin/bash
#SBATCH -J job_dep
```

```
#SBATCH -p compute
#SBATCH -n 1
#SBATCH -mem=100m
echo ${0}
sleep 120
$ cp job1.slurm job2.slurm
$ sbatch job1.slurm
Submitted batch job XXXXX
$ sbatch -dependency=afterok:XXXXX job2.slurm
```

5) running the same job (a.k.a same computation) on several parameters

10min: using a for loop in bash

10min: using SLURM job arrays

See: <https://groups.oist.jp/scs/running-matlab-jobs-sango-cluster>