

Using Hydra (IFT-cluster)

Chia-Hsun Chuang (Albert)

IFT-UAM/CSIC, Multidark Fellow

2013/11/12

Hydra

- Dual processor with 4 or 6 cores (8 or 12 cores/node)
- 70 nodes (704 cores in total)
- 24 GB DDR3 RAM



- To login:
`ssh -X username@hydra.ift.uam-csic.es`
- To submit a job: `qsub script.pbs`
- To observe the job sent: `qstat` or `showq`
- To delete a job: `qdel jobid`
- To test some code interactively, you might login to other nodes by `ssh hydraXX`,
where `XX=username*2`

- I will give you an "installed version" of cosmomc
"cosmomc_planck_hydra.tar.gz"
- In this tarball, you could also find an example script for submitting a job (cosmomc.pbs), the batch script which you should copy to your home directory and then re-login(.bashrc).
- You might compile or run some tiny jobs interactively, but **DO NOT** submit jobs yet, since our reservation of the cluster starts on Wednesday.