

# Using COSMA

A telescope for cosmology

Alastair Basden

Alastair Basden  
DiRAC / Durham University

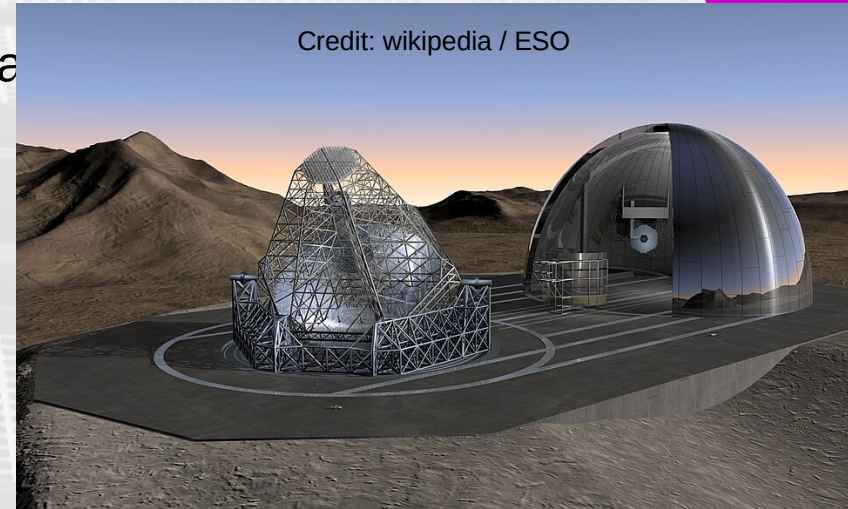


**DiRAC**  
High Performance  
Computing Facility



# Using telescopes

- Where are you pointing / where is your observation?
  - Below the horizon?
- What are the seeing conditions like?
  - And is it raining?
- What wavelength filters are in use?
- How long does the observation need to be?
- What is the sensitivity of the instrument?
- How do you use the data analysis platform / pipeline?
- You need to understand the physics behind what you're trying to achieve!



# Using cosmascopes

- You will benefit most if you try to understand a HPC system
  - They all vary slightly, the basics are the same
- If something isn't working try taking a scientific approach
  - Why, what has happened, how can it be resolved?
- And learn to make best use of it
  - Optimised codes, tuned to the system
  - Best for aiding net-zero
- If someone/website gives you instructions spend time to understand them
  - Don't just follow them blindly (sudo apt install just won't work!)

# Accessing the system

- First, sign up for an account on SAFE:
  - You'll need a password and authenticator app
- Upload an ssh key (the public part)
  - Protected by a passphrase
- Request a login account on COSMA
  - And join the do020 project
  - Once authorised, you'll get an email from COSMA
    - Read it!

# Logging in

- You will need:
  - Your ssh key (protected by a passphrase)
  - Your COSMA password
    - This will be sent in the welcome email
- ssh to COSMA may fail:
  - You've forgotten your passphrase or password
  - Your account has been disabled
    - Inactivity, failure to verify details on SAFE, invalid email
  - You are blocked temporarily
    - Too many password failures
  - You're using the wrong ssh key, or its file permissions are wrong
  - You're a Windows user!
- ssh -vvv will give you more information
- If you are asked to reset your password when you log in, please follow the instruction carefully

# Login nodes

- A shared resource
  - Think about what you run here
  - Don't use all the cores
  - Don't use all the RAM
  - Don't leave things running when not needed
- Good for:
  - Compiling code
  - Quick test jobs
  - Submitting to the queues
  - Managing files

# Slurm batch queues

- Slurm is your friend
  - Provides access to the compute resources
- Be inquisitive: learn the environment
  - What queues are available to you
  - How busy are they?
  - What type of nodes?
    - Cores, RAM, GPUs

# Storage

- Your home space:
  - `/cosma/home/do020/USERNAME`
- Application space:
  - `/cosma/apps/do020/USERNAME`
- Scratch space:
  - `/snap8/scratch/do020/USERNAME`
- Data space:
  - `/cosma8/data/`
- Each is optimal for different things – and can be tuned

# Storage usage

- HPC file systems typically like smaller numbers of larger files
  - Many small files can cause performance issues
  - Therefore think about how you will write your outputs
    - e.g. don't let your simulation produce 100k files per day!

# Software

- Available as “modules”
  - Set up your desired environment using
    - “module load ...”
  - Try to understand what the module command does
- Or install your own
  - Python virtual environments are useful for this

# Jupyter hub

- You'll need an ssh tunnel:
  - `ssh -L localhost:PORT1:node:PORT2 USERNAME@LOGIN.NODE`
  - Understand what this means (and check its output)
    - PORT1 on your laptop (localhost) will act like a wormhole
      - Transporting data
        - Via LOGIN.NODE
        - To/from PORT2 on node
- Close notebooks when you no longer need them

# Conclusions

- Learn to master your HPC systems
  - Don't let them master you!
- Visit the support pages
  - [cosma.rtf.d.ac.uk](http://cosma.rtf.d.ac.uk)
- Ask around for help and advice
  - As a last resort, contact [cosma-support@durham.ac.uk](mailto:cosma-support@durham.ac.uk)