A very quick introduction to HPC

(HPC = High Performance Computing)

Lucy Whalley // lucydot.github.io

Who am I?

- My name is Dr Lucy Whalley
- Vice-Chancellor's Fellow in Physics at Northumbria
- Member of the EE Faculty HPC steering group
- I use HPC for my research in solid-state physics

(HPC = High Performance Computing)

Quantum mechanical vibrations

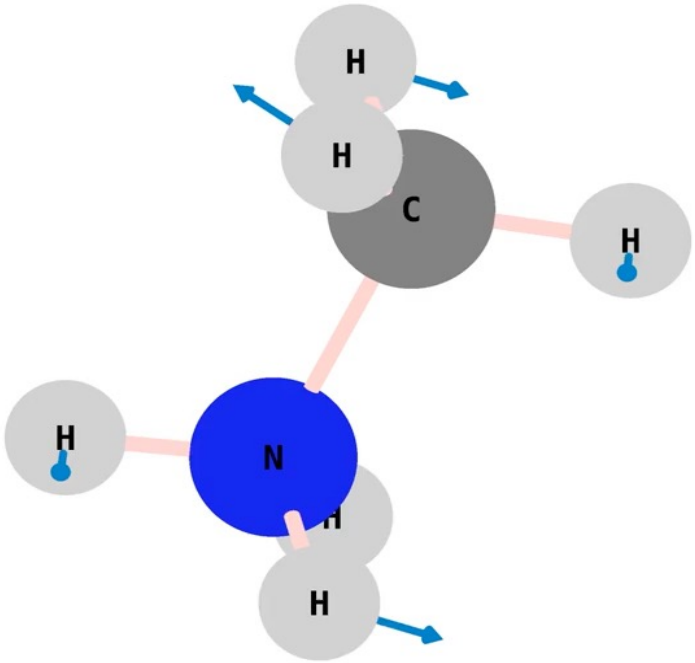

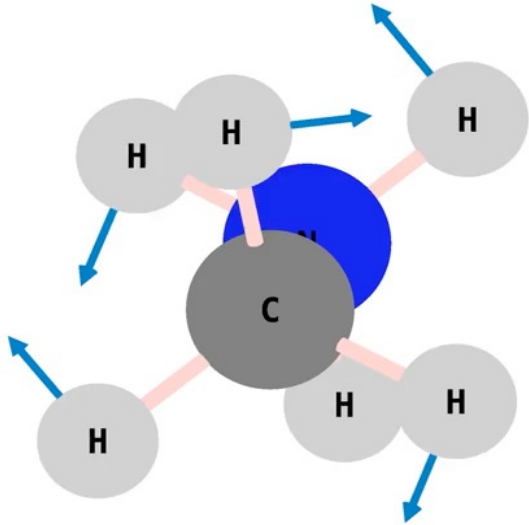

Courtesy Jarvist M. Frost

OUTLINE

01 Why use HPC?

02 HPC basics

03 Finding out more

04 Questions

Why use HPC?

Research problems can outgrow the computer where they started

- Statistical validation – running the same model 1000 times, each time takes an hour
- Increase in dataset size – the new data is 10 times as large and will crash a laptop
- Increase in model complexity – going from 2D to 3D simulations leads to a large increase in simulation time

Supercomputers (also known as clusters) are optimized for intensive computing

This calculation would take 24 days on my laptop (and would really slow down ~~Netflix~~ Teams meetings)

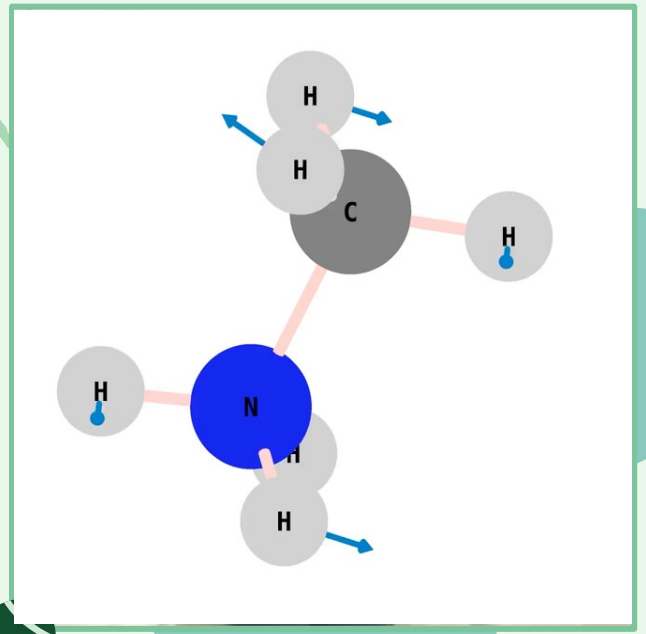

Why use HPC?

Using a cluster has several advantages

- **Speed** – clusters have *__lots of__* cores
- **Volume** – processing memory (RAM) and disk storage for large amounts of data
- **Efficiency** – as there are many system users
- **Cost** – bulk purchasing results in a reduction in cost
- **Convenience** – why tie up your own computer for hours if you can use another one?

02

HPC basics

What is a supercomputer?

What is a supercomputer?

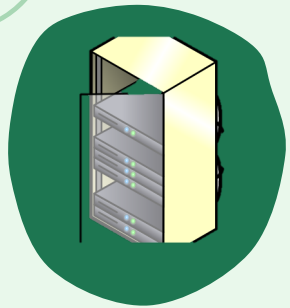

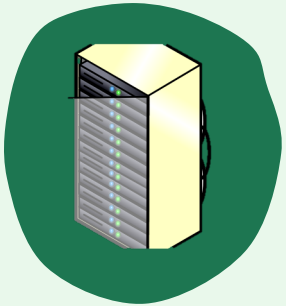

Laptop

2 cores

Server

16-128 cores

Supercomputer

~500 – 7.6 million cores

Top500.org

Rank	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	**Supercomputer Fugaku** - Supercomputer Fugaku, A64FX 48C 2.2GHz, Tofu interconnect D, Fujitsu RIKEN Center for Computational Science Japan	7,630,848	442,010.0	537,212.0	29,899
2	**Summit** - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM DOE/SC/Oak Ridge National Laboratory United States	2,414,592			
3	**Sierra** - IBM Power System AC922, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM / NVIDIA / Mellanox DOE/NNSA/LLNL United States	1,572,480			

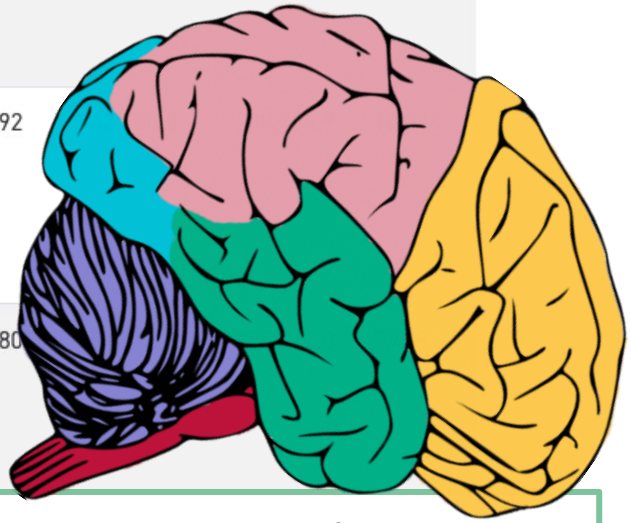

The path to exascale computing (10^{18} flops)

Which computers can I access?

HPC facilities are structured according to tiers

TIER	COMPUTER	Access route
0	Multiple	Prace - (http://prace-ri.eu)
1	Archer2	Multiple
2	Oswald	Raise an IT ticket

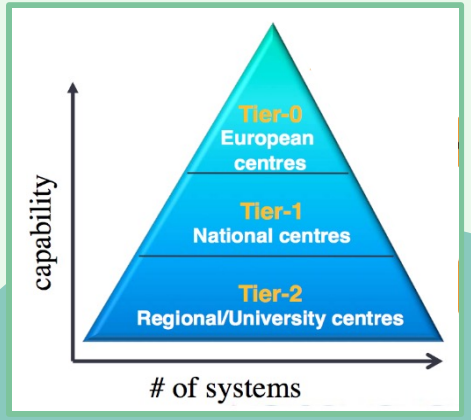

How do I access the computer?

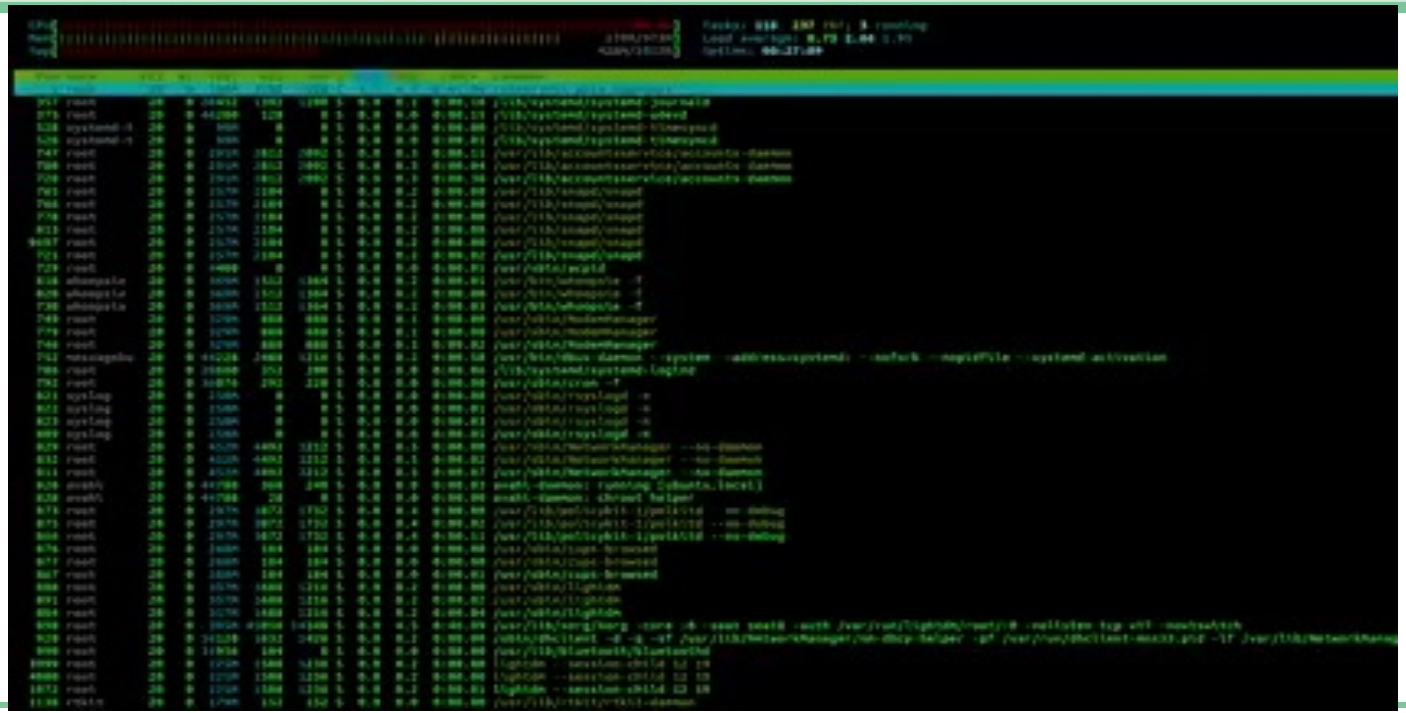

Using a terminal (just like the movies)

How do I access the computer?

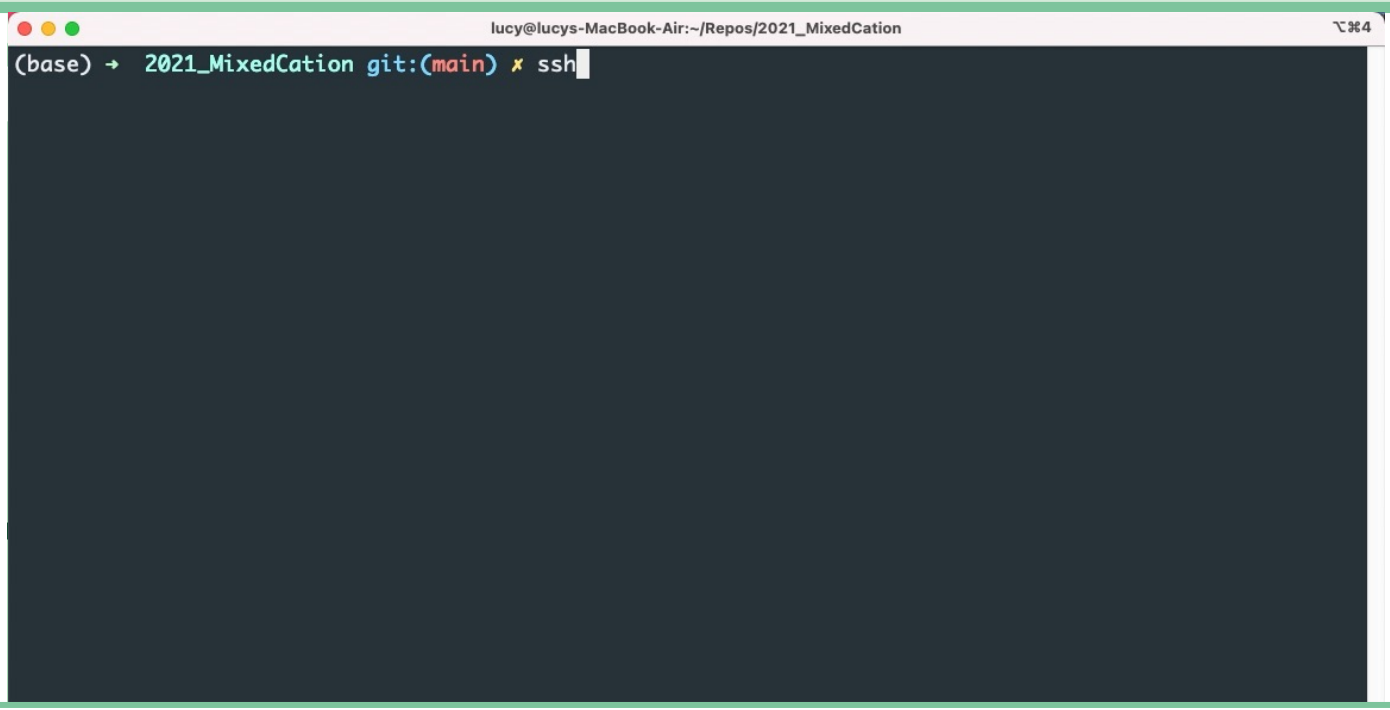

Using a terminal and Unix commands

How are resources distributed?

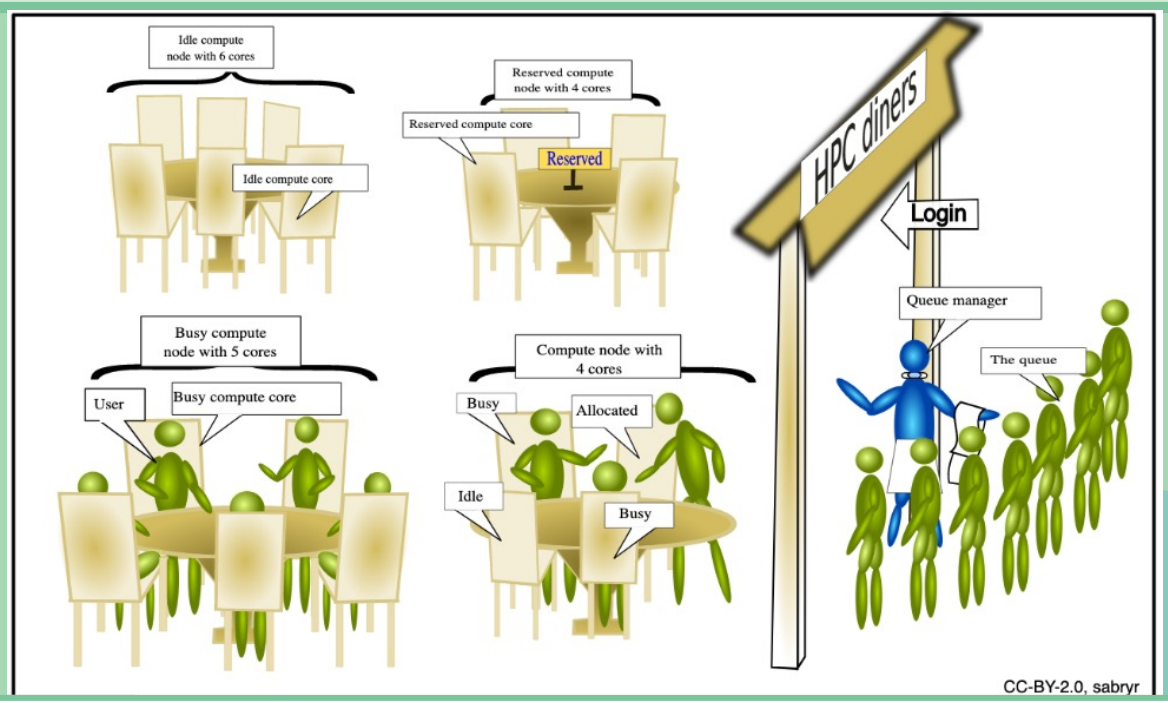

A schedular is used to manage the job queue

03

Finding out more

What next?

Where can I get more support?

Departmental HPC rep

MCE: Craig Warren

RSE community rep

Lucy Whalley

To access Oswald

Raise an IT Ticket

Archer2 national cluster

www.archer2.ac.uk

HPC workshops

www.archer2.ac.uk/training/

PRACE European cluster

https://prace-ri.eu/

THANKS!

slides: lucydot.github.io/talks

Watch out for the Research Software Engineering @ Northumbria event this autumn

Credits: Slidesgo, Flaticon, Freepik, epcced.github.io/hpc-intro-2020-06-30/