I Just Can't Use My Laptop Anymore!

What can I use instead?

Motivation

- Experience working with researchers tells me that it's not always obvious what the routes to solving this problem are
- Sometimes it isn't even obvious that there is a problem at all!
- UoY offers a range of computing services to researchers, can be difficult to know what is appropriate
- Learning curve for new things can be demotivating / kill momentum / not fit to timescales
- I want to see how many live demos of different things I can fit into one talk

Objectives

- Identify key reasons why that trusty laptop isn’t up to the challenge
- Signpost UoY resources that are available to you, which might help to solve the problem
- Touch on how you might use some of these resources to solve problems
- Field any immediate questions you have about laptop destroying software you are using

Disclaimer

I know the answers will just seem like USE BIGGER COMPUTER!

But *which* bigger computer should you use? And *how*?

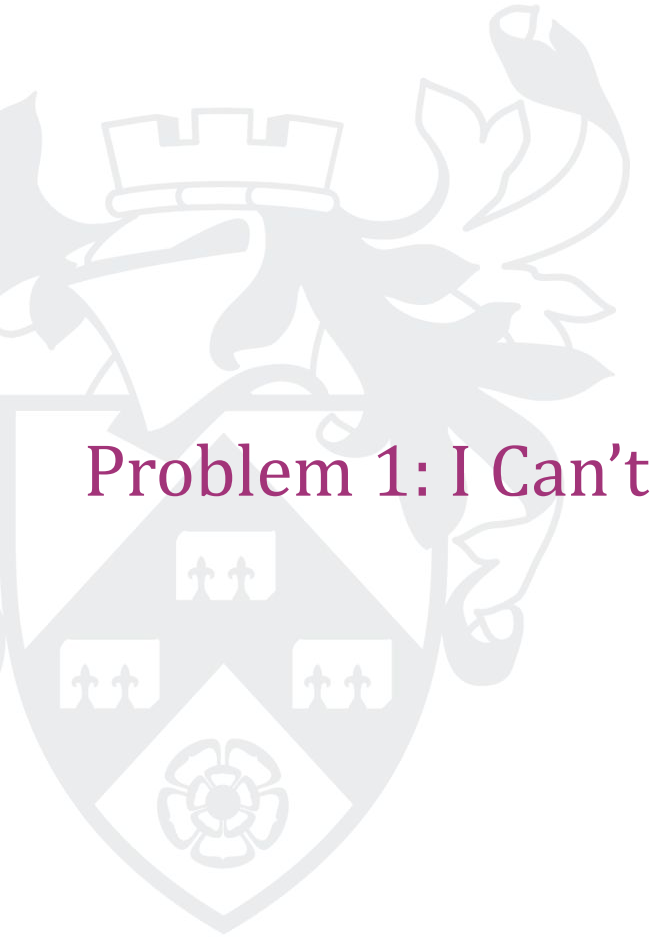

Problem 1: I Can’t Use My Laptop Whilst It’s Running!

Why Does This Happen?

- Software is resource-intensive, preventing other things from running at a usable level
- Incredibly frustrating to have your one computer grind to a halt whenever you run something intensive
- The lowest effort solution to this problem is to find another computer on which you can run your software

Where Can I Find Another Computer?

Option 1: Virtual Desktop Service (VDS)

- Quick and easy access to a University of York Windows 10 desktop
- Familiar software and files available by default
- Not so suitable for Big Software

Live Demo 1/N

Where Can I Find Another Computer?

Option 2: Interactive Linux Service

- Comparable to VDS, but for Linux
- Fully-featured Linux desktop environment (matches campus machines)
- Familiar software and files available by default
- Can be a little competitive

Live Demo 2/N

Problem 2: It Crashes When I Run It With More Data!

Why Does This Happen?

- Need to re-run software with a larger dataset in response to reviewer comments
- Found an exciting new dataset that you want to explore
- Software is not written to make efficient use of memory
- The software just needs a lot of memory to run beyond a certain problem size

How Can I Check This?

- Windows: the ‘Performance’ tab of the Task Manager can show you system memory usage
- MacOS: the ‘Memory’ tab of the Activity Monitor can show you system memory usage
- Linux: your distro’s System Monitor can show you system memory usage

Live Demo 3, 4, 5 / N

Where Can I Find More Memory?

UNIVERSITY of York

Option 1: Interactive Linux Service:

- 2TB of memory across 4 nodes!
- Learning curve is gentle - directory navigation, loading software, running software
- Introductory tutorial available on the wiki

Live Demo 6/N

Where Can I Find More Memory?

Option 2: Viking

- Standard Viking node has 192GB of memory (170 of these)
- High memory nodes have 768 - 1.5TB of memory (6 of these)
- Completely free!

Where Can I Find More Memory?

Option 2: Viking

- Learning curve a little steeper than Interactive Linux Service
- Graphical desktop available if desired
- Interactive sessions available if desired
- Work managed by a workload manager, Slurm
- Wide range of software available as modules

Live Demo 7/N

Problem 3: It Takes Too Long!

Why Does This Happen?

- Software is often not designed with performance in mind
- Key algorithms might not scale well with increasing problem size
- Software may not exploit potential parallelism
- Hardware might be old or underperforming
- The task might just take a long time

Can I Make My Software Go Fast?

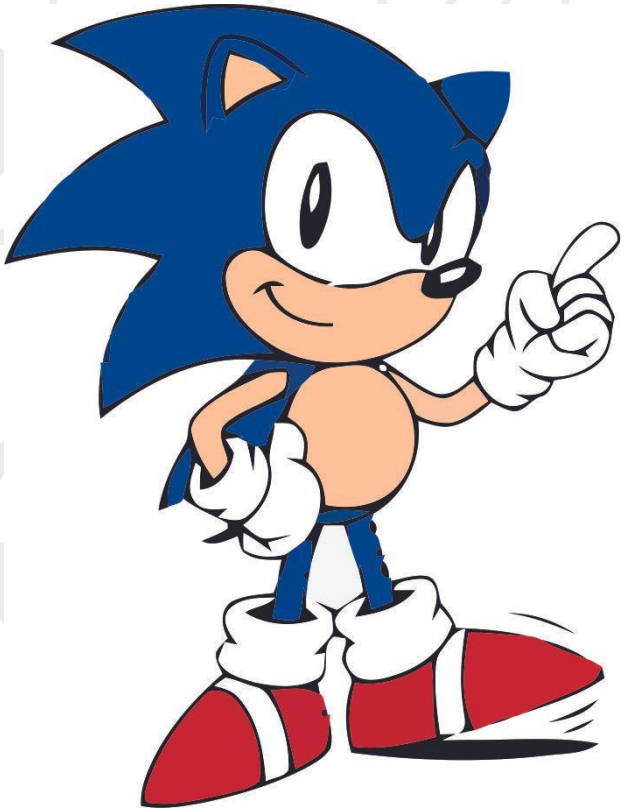

- You can certainly try!
- Considering performance is a good habit to get into - computation ~= energy usage
- Report suspected performance issues to developers, if the code is not yours
- Try profiling your code! (Another talk…)

Where Can I Run For Long Time?

Option 1: Viking

- Up to 48h, 7 day, and 30 day jobs are allowed to run
- Requires basic understanding of the Viking workload manager
- Documentation available online
- Help available from IT Services
- Completely free!

Live Demo 8/N

Problem 4: It Uses Too Much Disk Space!

Why Does This Happen?

- New developments in your field produce larger datasets
- Software might need large amounts of input data to run at all
- You might want to generate more output data for a more detailed analysis
- You might be working on several things at once

Where Can I Get More Storage?

Option 1: Departmental

- Academic departments typically provide some storage to researchers
- This is backed up, and suitable for storing work you don't want to lose whilst a project is ongoing
- Usually available on VDS, Interactive Linux Service, SFTP etc.
- Contact your Departmental Computing Officer (DCO) for info

Live Demo 9/N

Where Can I Get More Storage?

Option 2: Viking

- Viking contains a 2.5PB file system for storing data that supports your computational work on Viking
- There are ***NO BACKUPS*** - if you accidentally delete files, they are **gone**
- Use in combination with option 1!

Live Demo 10/N

Where Can I Get More Storage?

Option 3: Vault

- IT Services provide archival storage to researchers
- Suitable for large volumes of data that needs to be accessed *very* infrequently
- Up to 500TB free!
- Use in combination with options 1 & 2

Live Demo 11/N

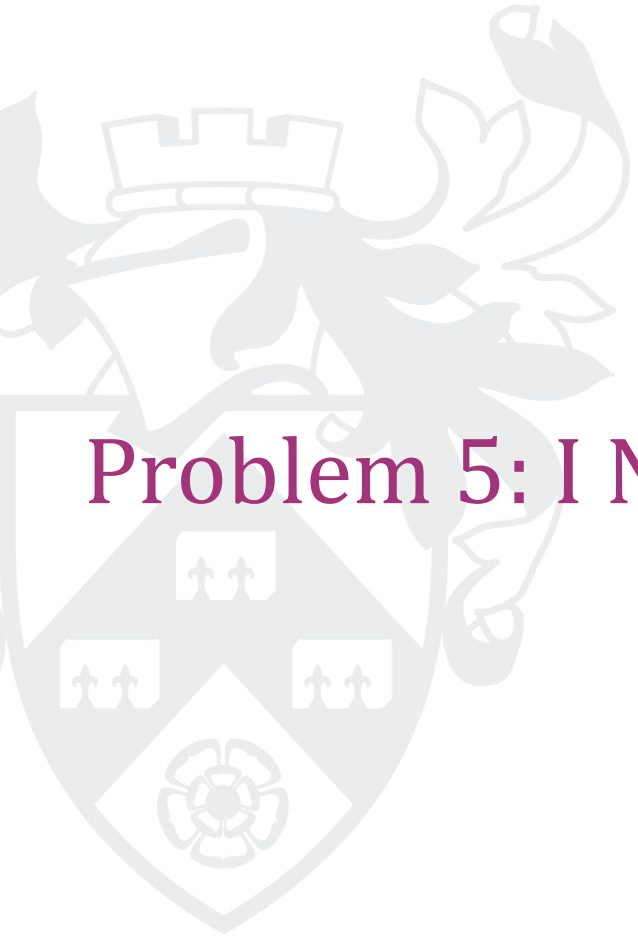

Problem 5: I Need a GPU!

Why Does This Happen?

- GPUs are becoming increasingly prevalent in research software
- Some techniques benefit greatly from their usage
- Some software packages that you want to try are developed to run only on GPUs
- Many people are interested in exploring the potential applications of GPUs

Where Can I Find A GPU?

Option 1: Viking

- Viking has 2 GPU nodes, each containing 4 x Nvidia V100 (read *powerful*) GPUs
- As straightforward to use as anything else in Viking
- Not suitable for small-scale / short running GPU work

Live Demo 12/N

Where Can I Find A GPU?

Option 2: Bede

- GPU-first regional computing facility
- ~128 x Nvidia V100 (again, *powerful*) GPUs, 16 x Nvidia T4 (ML inference) GPUs
- Free access via your project supervisor
- Non-standard processors, so not everything will work straight away! IT Services can help

Live Demo 13/13

Conclusions

- A wide range of resources available to you, almost all for free
- The resources can all work together to help you to solve problems
- Documentation is available in the UoY Wiki
- IT Services can help - itsupport@york.ac.uk
- Coding Club can help - come to a drop-in session!

Useful Links

- Virtual Desktop Service (VDS)
- Interactive Linux Service
- Viking
- DCO list for departmental storage queries
- File transfer (SFTP & SCP Services)
- N8 Bede
- Data Safe Haven
- <itsupport@york.ac.uk> - if you need support with any of this!