Introduction to Advanced Research Computing

The ARC Team

https://www.arc.ox.ac.uk

https://www.it.ox.ac.uk

Outline

Introduction

Who we are...

The Advanced Research Computing (ARC) service is a central resource available to any Oxford University researcher who needs access to research computing services from any Division or Department.

While we have a range of high performance computing (HPC) systems available, ARC offers much more than just hardware.

ARC users have access to full time support staff specialising in HPC hardware and software. Support is available to users throughout their research project.

High Performance Computing

There is no single, all-encompassing definition for HPC.

For our purposes it is:

> *Computing which cannot be achieved in a reasonable time on a laptop or desktop workstation and may be carried out on multiple processors.*
>
> ***HPC is the ability to do work in less time, do more work in the same time, or achieve something that is otherwise impossible.***

HPC does not always mean running code in parallel on a cluster — access to a single large memory ‘fat’ node — is just as valid.

High Performance Computing

But how fast is fast enough?

- A desktop PC can deliver tens of Gflops (flops = floating point operations per second). Tens of millions is a lot of flops but that may or may not be enough...
- An extreme example: short range weather forecast: prediction for the next day must be delivered in less than 1 day;
- The Met Office models translate into a requirement for ~1 Pflops, that is 1 million times more that the PC, so the simulation has to run on many CPUs, working together on the same model.

The fastest supercomputers achieve Petaflops/Exaflop performance.

Parallel computing is for all

For the past two decades, the trend in supercomputer design has been towards:

- **clusters** (compute servers connected by a fast local network) and
- **commoditisation** (off-the-shelf components, economies of scale)
  - x86 is the dominant instruction set (Intel and AMD CPUs)
  - Linux is the dominant OS

At the same time, CPU design incorporates more and more cores.
The difference between your departmental machine and a supercomputer is limited to:

- server grade hardware (not much faster but more reliable, maybe more RAM)
- fast (usually Infiniband) network and fast parallel storage

HPC (traditionally the preserve of high-end computing in physics and engineering) has become a challenge for everyone (software vendors, researchers, etc.) everywhere

- **supercomputers** as well as
- **servers** and **workstations** (with multicore CPUs)

Who needs HPC/HTC?

<table>
<tr><td>Researcher 1</td><td>• Has a large volume of independent jobs (e.g. processing video files, genome sequencing, parametric studies)
• uses serial applications
• needs to accelerate research</td><td rowspan="2">High Throughput Computing (HTC)
(many computers):
• large number of jobs
• one job runs on one cluster node
• one job can harness the 48 cores available per node — this is achieved through multi-threading of via coordinating concurrent processes</td><td></td></tr>
<tr><td>Researcher 2</td><td>• developed serial code and validated it on small problems
• to publish, needs some ‘big problem’ results
• hits a performance wall</td><td rowspan="2">High Performance Computing (HPC)
(single parallel computer):
• static environment
• single large scale problems
• tightly coupled parallelism</td></tr>
<tr><td>Researcher 3</td><td>• needs to run large parallel simulations fast (e.g. structural mechanics, molecular dynamics, computational fluid dynamics)</td><td></td></tr>
</table>

How can ARC help you?

Researcher 1	• Has a large volume of independent jobs (e.g. processing video files, genome sequencing, parametric studies) • uses serial applications • needs to accelerate research	**HTC:** • large number of jobs • one job runs on one cluster node • one job can harness the 48 cores available per node — this is achieved through multi-threading of via coordinating concurrent processes
Researcher 2	• developed serial code and validated it on small problems • to publish, needs some 'big problem' results • hits a performance wall	**Parallel code development** • parallel execution may be a solution • this course can be a start • seek ARC advice
Researcher 3	• needs to run large parallel simulation fast (e.g. structural mechanics, molecular dynamics, computational fluid dynamics	**HPC** • single large scale problems • jobs run on multiple nodes • many commercial and free applications can be run across multiple nodes

Examples of research using ARC

- Anthropology
- Chemistry
- Computational Medicine
- Engineering
- Novel Architectures: SKA
- Physics
- Quantum Technology
- Sociology
- Statistics
- Zoology
- Philosophy and Oriental Studies

What does ARC offer?

- Central University HPC;
- User Support and Training;
- Access to range of systems and technologies;
- Premium Services;
- Support for Commercial Customers.

ARC HPC Services

HPC Clusters

Cluster services:

- HTC - high throughput
- ARC - large scape parallel

GPU nodes on HTC include:
Pascal, Volta, Turing, Ampere, Grace Hopper, MI250X

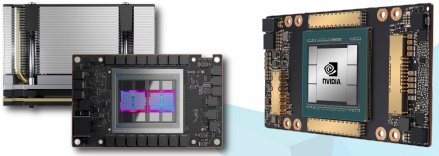

Hosted at Begbroke Data Centre

HTC Cluster

- Minimum job size: 1 core
- CPU and GPU node on this cluster include:
  - High Memory Nodes
  - Single Precision GPU Compute Nodes
  - Double Precision GPU Compute Nodes

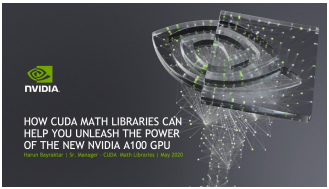

https://arc-user-guide.readthedocs.io/en/latest/arc-systems.html

ARC Cluster

ARC provides a capability cluster comprising of:

- >250 general compute nodes; designed for multi-node parallel computation
- OS: CentOS 8.x (two nodes for legacy software running CentOS 7.7)
- Scheduler: SLURM

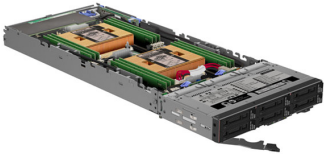

>14 000 CPU cores

All connected with fast low-latency network/interconnect (InfiniBand)

https://arc-user-guide.readthedocs.io/en/latest/arc-systems.html

Storage

The ARC clusters are connected to specialised high performance storage.

Storage performance is extremely important for HPC systems — and is therefore expensive.

ARC offers 3 Petabytes of usable high performance OnTap storage, as well as 1 Petabyte of ultrafast, low latency, NVMe SSD WEKA scratch storage.

https://arc-user-guide.readthedocs.io/en/latest/arc-storage.html

Linux Operating System

All ARC systems operate on the Linux Operating System, rather than on platforms like Windows or MacOS.

Why Linux is the Preferred Choice for HPC:

- Cost-Efficiency
- Dependability
- Optimal Performance

To use ARC, a basic knowledge of Linux and the command line interface is essential. Numerous online resources are available to help you acquire this understanding.

See the ARC training pages for more information:
https://www.arc.ox.ac.uk/training/

User Support

User Support

- ARC Training Courses
  - Getting Started with the ARC clusters
  - Effective Use of Clusters (for non-programmers)
- User Documentation
- Attend student welcome events
- Advise on costing grants
- Installing, testing software applications
- Advice on using software, best practice and examples

https://arc-user-guide.readthedocs.io/

Software Support

- Provide application software support across a broad range of research disciplines;
- Install new or updating application versions upon request;
- Provide build and installation assistance to users with their own (or highly customised) applications;
- Receive on average 4 application requests per week.

https://arc-software-guide.readthedocs.io/

Available software

https://arc-module-list.readthedocs.io/

ARC Accounting

ARC uses a ‘credit’ system to keep track of resource usage on the ARC clusters, with one credit translating to running for one second of wall-clock time on a single CPU core.

ARC usage credits are ‘free’ for researchers in the University of Oxford, in the sense that researchers do not have to pay directly to gain access to the ARC service.

However purchasing credits (e.g. from research grants) gives access to an elevated level of service, the main benefit of which is a more rapid turnaround of jobs.

All new ARC projects are set up with an allocation of standard or basic compute credit.

Further credits can be obtained upon request see:
https://www.arc.ox.ac.uk/arc-accounting/#RequestCredit
https://www.arc.ox.ac.uk/arc-service-level-agreements/

Getting started with a user account...

All ARC users must be a member of an ARC Project. You can either:

- Join an existing project:
  The project owner will be asked to authorise your application.
  https://www.arc.ox.ac.uk/arc-user-registration-page/
- Request a new project:
  (you must be a member of university staff)
  Once created you can request a user account.
  https://www.arc.ox.ac.uk/arc-project-registration/

Training

Further Training

- **Introduction to Linux**
  external training for essential Linux command line knowledge
  https://www.netacad.com/courses/linux-unhatched/
  https://ryanstutorials.net/linuxtutorial/
- **Introduction to ARC** (this one!)
- **Getting Started with ARC**
  Online training for new ARC users
  https://www.arc.ox.ac.uk/getting-started-arc/
- **Effective Use of HPC Clusters**
  Advanced online training of more experienced users
  https://www.arc.ox.ac.uk/effective-use-hpc-clusters/

https://www.arc.ox.ac.uk/training/