



Introduction - Advanced Research Computing

Introduction to ARC

- Who we are?
- What do we offer?
- What is High performance Computing?
- Where do I start ?

Who we are ...

- Advanced Research Computing (ARC) is a central resource available to any Oxford University researcher who needs access to research computing services from any Division or Department.
- ARC users have access to full time support staff specialising in HPC hardware and software. Support is available to users throughout their research project.
- User Support and Training
- Access to range of systems and technologies
- Premium Services
- Support for Commercial Customers
- Courses on topics such as: programming, code optimisation, debugging

What do we offer



Provide access and Support for hardware and software

Parallel Programming Made Easier

Standard C Code

```
void saxpy_serial(int n,
                  float a,
                  float *x,
                  float *y)
{
    for (int i = 0; i < n; ++i)
        y[i] = a*x[i] + y[i];
}

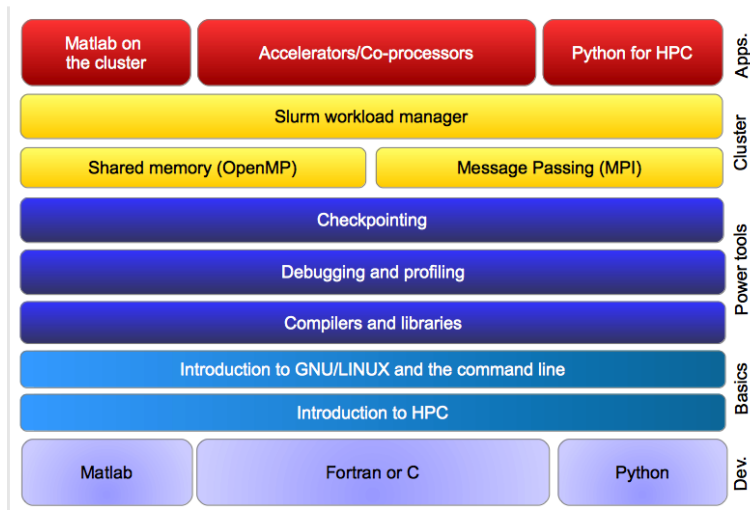
// Perform SAXPY on 1M elements
saxpy_serial(4096*256, 2.0, x, y);
```

Parallel C Code

```
__global__
void saxpy_parallel(int n,
                    float a,
                    float *x,
                    float *y)
{
    int i = blockIdx.x*blockDim.x +
            threadIdx.x;
    if (i < n) y[i] = a*x[i] + y[i];
}

// Perform SAXPY on 1M elements
saxpy_parallel<<<<4096,256>>>>(n,2.0,x,y);
```

Software installation and general support

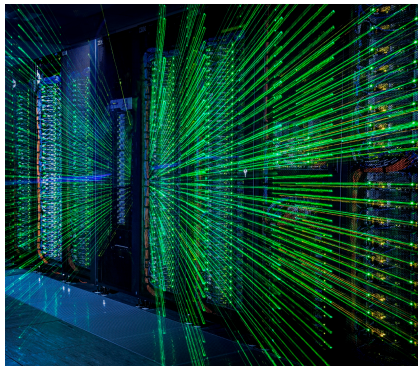


Software installation and advice



Data Centre in Begbroke - IAT

High Performance Computing



There is no single, all-encompassing definition for HPC For the purposes of this course it is computing which : cannot be done on a desktop or a server 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.**

Linux Operating System

All ARC systems run the Linux Operating System (as opposed to Windows, MacOS). Linux is commonly used in HPC because:

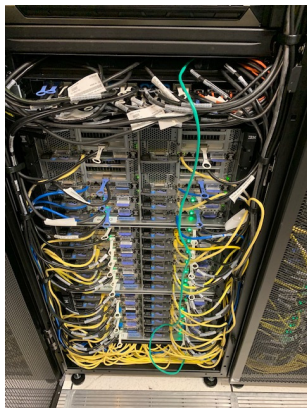
- Cost
- Reliability
- Performance

Linux is also used in:

- • Non-HPC servers
- • Android phones
- • Smart-devices

The Linux logo, which consists of the word "Linux" in a green, monospace-style font, followed by an underscore character "_". It is set against a solid black rectangular background.

HPC clusters



Cluster services: Access and Support for hardware and software

- arcus-htc – high throughput
- arcus-b – large scale parallel

Graphical Process Units

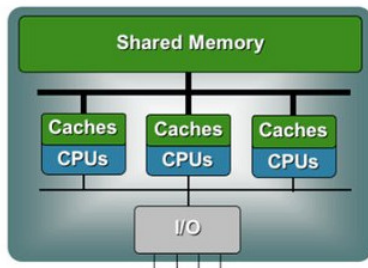


ARC GPUs include: Intel Xeon PHI, K80, DGX1V, DGXQMAXQ

- The ARC clusters are connected to specialised high performance storage.
- Storage performance is extremely important for HPC systems – and is therefore expensive.
- A new storage system was installed in 2019. The total usable storage capacity is currently two Petabytes (2000 terabytes) of usable high performance storage.

Shared Memory Systems

- Lots of processors and memory in a box
- All sub processes in one box
- Acts like a large ‘desktop PC’
- Applications split into “threads”
- Usually programmed in something like OpenMP
- Easy to program
- Financially expensive at large scale
- Good for processes with common data e.g. stats, genomic



arcus-htc is a High Throughput Cluster to support users with lower-core count workloads. OS: CentOS 7 Scheduler:SLURM.

- Single core jobs, and sub single node core count jobs account for a substantial proportion of the jobs submitted by users to both Arcus systems.
- arcus-htc supports users with lower-core count workloads and will automatically pack users jobs to share compute nodes, thus increasing job throughput.
- The HTC cluster allows non-exclusive use of nodes.

CPU and GPU nodes on this cluster include:

- 104 Intel Sandybridge 16core 64GB (memory per node)
- 4 Intel Ivybridge 16core 128GB (memory per node)
- 5 Intel Broadwell 40core 512GB (memory per node) 8XV100-LS-8GPUs per node NVidia DGX-MaxQ
- 1 Intel Broadwell 40core 512GB (memory per node) V100 8GPUs per node NVidia DGX1V
- 3 Intel Knights Landing 288core 192GB (memory per node) Xeon Phi
- 3 Intel High Memory Nodes, 2 with 1.5TB memory 32 cores per node), 1 with 6TB (128 cores)

Graphical Processing Unit

- Details of GPU nodes on arcus-htc cluster
<http://help.it.ox.ac.uk/arc/using-gpus>
- Efficient, good for highly parallel jobs
- High memory bandwidth
- Significant speedup for some problems

Cluster (Distributed memory)

- Problem split across smaller machines with software, typically MPI
- Sub-tasks occasionally exchange information
- Cheap to build
- Hard to program
- Good for partially self contained problems e.g. fluid flow
- All connected with fast/interconnect (QDR Infiniband)




arcus-b is the system designed for multi-node parallel computation.

OS: Centos 6.6, Scheduler: Slurm

- IBM/Lenovo System
- Around 400 compute nodes
- Each standard node has : 16 compute cores (Haswell) or 20 Compute cores(Broadwell) , 64, 128 or 256 GB RAM

Do I need to be a computer specialist?



case study

Collisions and Complexes of Free Radicals

Sarantis Marinakis from the Department of Chemistry used the ARC facilities for his project: *Collisions and Complexes of Free Radicals*. Details about this research can be found on the ARC website. Below he responds to questions about his experience of using the ARC facilities.

What were the main reasons you started using the facilities at ARC?

My programmes sometimes need more than 44 Gbytes, so it was the access to a very large amount of RAM memory that was key for me. I was relieved to discover there were computing facilities on that side within the University. And that there's such a lot of support to get your programmes running.

What kind of support have you received from the team at ARC?

I was given a lot of help compiling a very tricky programme – H302on. The ARC team gave me liveable advice in one-to-one meetings whenever I needed it.

What other benefits have you experienced from using the ARC facilities?

So much that it was very useful was that I was able – with the help of the team – to try out programmes that were not already installed. It was also fantastic that a few software companies allowed us to try out some expensive programmes at ARC for free.

What would you say to researchers who are thinking of using the facilities at ARC?

I'd recommend they contact the team and discuss their plans. Some people may be put off if they can't use the software they want to use on the list of those available, but ARC will get hold of new software if there's sufficient demand for it and it's not too expensive, so the list is continuously expanding in response to the needs of users.

case study

Mapping the Muslim population of Uttar Pradesh

Raphael Susewind is an Associate of the Contemporary South Asian Studies Programme at Oxford while he completes his PhD research into the politics and poetics of Muslim belonging.

I first started using the facilities of ARC when I was in India working on my PhD research and wanted to create a map showing the distribution of the Muslim population in Lucknow. Since such data is not officially available, I had to make informed guesses based on the religious correlations of names listed in each neighbourhood's public election register. I asked myself whether – rather than just doing this for fun – I could scale up my algorithm to the whole state of Uttar Pradesh, which has a population of close to 200 million. I am excited that my desktop would take about four months for this task – and most likely longer because something would be bound to go wrong at some point. So I contacted the team at ARC – with their resources I could do the work in less than 10 days.

I got my algorithm running, but I was wondering why it wasn't processing as quickly as I had expected when I received an email from ARC pointing out I was using 1 GB per node rather than 16. I had simply asked the ARC high-performance computer to do the same as my desktop, rather than exploiting the power of parallel computing.

The team helped me rethink my approach and things speeded up considerably after that. I'm so grateful that they decided up on how my work was progressing – it's this kind of attention to supporting users that I value most about ARC. The computers are not simply very big desktops, because they handle masses of data, they offer the potential for researchers to think in a different – and much more imaginative – way about the work they want to undertake.

Supervised machine learning

Cat vs dog

Recipe:

- Collect **labelled** data
- Build fancy model
- Train model by correcting its mistakes
- ...deploy

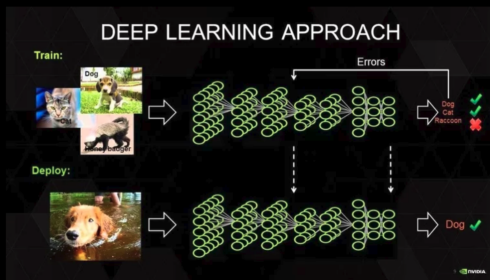
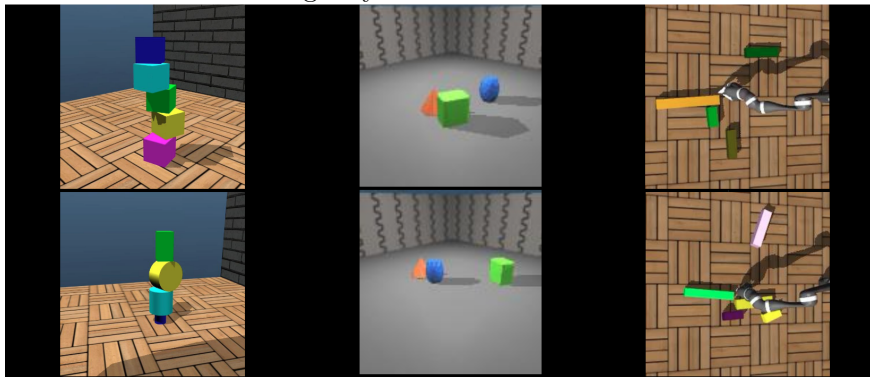


Image from: <https://www.datascience.com/blog/beyond-deep-learning-a-case-study-in-sports-analytics>

Learning Physics without text books



Examples of Learned Intuitions

instability detection, affordance estimation, visual forecasting



Constructive quantum interference in a bis-copper six-porphyrin nanoring. Sabine Richert and co-workers report on the exchange interaction, J , between two spin centres is a convenient measure of through bond electronic communication.



Please acknowledge ARC using doi : [10.5281/zenodo.22558](https://doi.org/10.5281/zenodo.22558)



Compilers + Libraries

Intel C/C++
Intel Fortran
Intel MKL
Anaconda
Python
Java
BLACS
HDF5
ParMetis
PETSc
Scalapack
VTK
MUMPS
NetCDF



Genomics

AutoDock
GMOD
BLAST
CUDASW++
HMMER
Mmseqs2
MUSCLE
MUMmer
Pymol
Bowtie
Nonviolent
Stampy



Chemistry

ADF
Gaussian03
Gaussian09
CPMD
SIESTA
CP2K
GAMESS-US
GROMACS
NAMD



Abaqus -FE
CFD-ACE
Stata
R
Rmpi
Matlab
Mathematica

This is just a subset of the more popular applications installed. The command : 'module avail' on ARC Services will show applications and versions.

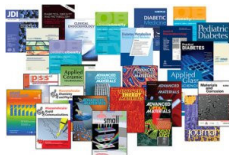
In Numbers



Over 40000 jobs per month



Over 500 new users per year

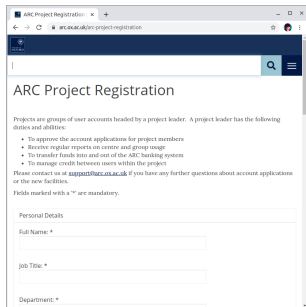


Minimum of 1 Publication per week referencing ARC



Projects worth over 50M pounds of research using ARC

How do I get on the System



The screenshot shows a web browser window with the title 'ARC Project Registration'. The address bar shows the URL 'arc.ac.uk/arc-project-registration'. The page has a dark blue header with a search icon and a menu icon. The main content area is white and contains the title 'ARC Project Registration'. Below the title, there is a paragraph explaining that projects are groups of user accounts headed by a project leader, followed by a bulleted list of duties and abilities. A contact email 'support@arc.ac.uk' is provided. A note indicates that fields marked with an asterisk are mandatory. The form section is titled 'Personal Details' and contains three input fields: 'Full Name: *', 'Job Title: *', and 'Department: *'.

ARC Project Registration

Projects are groups of user accounts headed by a project leader. A project leader has the following duties and abilities:

- To approve the account applications for project members
- Receive regular reports on centre and group usage
- To transfer funds into and out of the ARC banking system
- To manage credit between users within the project

Please contact us at support@arc.ac.uk if you have any further questions about account applications or the new facilities.

Fields marked with a "*" are mandatory.

Personal Details

Full Name: *

Job Title: *

Department: *

- Each user has an individual account associated with Arc.
- If you are associated with more than one project registered on ARC you will have a single id but two DATA directories, one for each of the projects.

Familiarity with Linux

You will need a little Linux knowledge to do this course

This course will not teach you Linux as there are other courses that can do that.

However you do not need an in-depth knowledge of Linux to: Do this course or to use HPC clusters including ARC systems



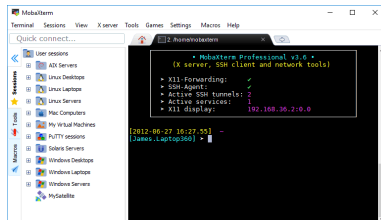
ARC provides users with two storage areas on high performance storage.

- HOME /home/username
Small quota. Used during login Home areas : 15GB per user
- Shared Project Space
DATA /data/project/username/
5TB per group
Large quota shared between project members
- SCRATCH In order to get the most efficient storage performance
ARC users can use our scratch storage for their compute jobs.
This storage will be automatically purged of data that has not
been touched for over two weeks

- **ARC makes best effort to ensure the integrity of data stored on our facilities** However, we are under no obligation to guarantee the integrity or availability of data this is the responsibility of the individual user.
- **No Backups limited snapshots of home and data** ARC does not accept any liability, financial or otherwise for loss of data. We recommend that users employ standard industry practice for their important data and store it at sites other than the ARC , for example, in their department.

Logging onto ARC Systems

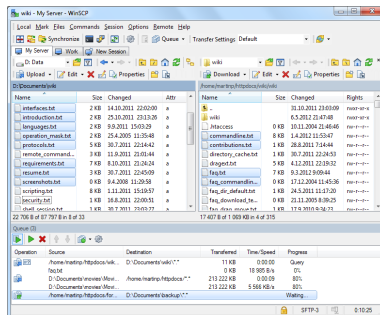
Remote Access Linux, Mac (and other Unix/Unix-like) users should use ssh (secure shell) to connect to the ARC systems from a terminal `ssh -X username@systemname.arc.ox.ac.uk` Windows users should download and install MobaXterm or PuTTY (ssh client for windows) and Xming for X11 support



Batch Jobs Jobs are submitted to the scheduler using the “sbatch” command and a job submission script. Never run intensive jobs on ARC systems without using the job schedule

Copying data to/from ARC

You can use either scp (secure copy) or secure file transfer protocol (sftp) MacOS (terminal) and Linux have scp and sftp built in. Windows users can use Mobaxterm or WinSCP



User Password Management

- User account management is performed on `myaccount.arc.ox.ac.uk` (use `ssh` to connect)
- Passwords can be changed by running the “`passwd`” command.
Enter login(LDAP) password:
Enter new passwd:
Re-enter new passwd:

Forwarding email addresses can be changed, please ask the ARC team.

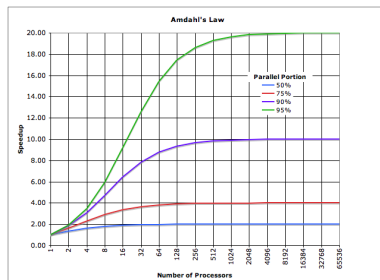
The accounting system explained:

- Facilitate charging / monitoring of jobs and usage of credit accounts
- All jobs (regardless of whether they are free or charged) consume credits
- Credits are usually consumed at the project level
- Use the command mybalance to check your credits

Common problem on arcus-b cluster: “**JobHeldAdmin**” = normally means your project has run out of credit

Using the right resource

Does your job require GPU nodes? Consult us Can you application benefit from multiple compute nodes? Splitting the task up isn't always straightforward. Remember you need to move your data to ARC storage ! Not a magic bullet approach – speedup and scaling Communication much slower than computation.



Submit all technical queries using the email address :support@arc.ox.ac.uk Make sure you include the following information to speed up response :

- Your user id
- screen shots if you are reporting login issues
- Job number if you are enquiring about an aborted/failed job
- the script used to submit the job or name of the script as well as its location

Our website www.arc.ox.ac.uk

User Guides <http://help.it.ox.ac.uk/arc/index>

For more information on ARC email theteam@arc.ox.ac.uk

For technical support email support@arc.ox.ac.uk