

Introduction to High Performance Computing for Life Scientists

Partners



Funding



Reusing this material



This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License.

http://creativecommons.org/licenses/by-nc-sa/4.0/deed.en_US

This means you are free to copy and redistribute the material and adapt and build on the material under the following terms: You must give appropriate credit, provide a link to the license and indicate if changes were made. If you adapt or build on the material you must distribute your work under the same license as the original.

Course Parameters

- Prerequisites
 - This course is designed to allow any researcher from the computational life sciences to be able to participate in and complete, regardless of their prior experience of high-performance computing.
 - Previous familiarity with the Linux command line is useful, but not assumed, and guidance is provided
 - No knowledge of programming is required
- Hands-on practicals form an integral part of the course
 - Use the UK national HPC service ARCHER
 - Learn by doing, gain practical skills and insights
 - Demonstrators will help with these

Aims

- What is HPC?
- Why do people use HPC and what do they use it for?
- Understand computer hardware
 - Which parts matter for performance in scientific applications?
- Understand processes and threads
 - How applications run on hardware
- Understand parallel programming models
 - How applications tackle problems in parallel
- Gain experience using an HPC machine
 - Dealing with common stumbling blocks

Aims

- Know how to evaluate parallel performance of an application
 - How do you know whether you're making good use of HPC resources?
- Understand current HPC architectures
- Know about parallel programming libraries
- Appreciate some of the challenges running life science pipelines / workflows on HPC systems
- Know about the UK & EU HPC landscape
- Gain an appreciation of the future of HPC
- Understand how HPC can benefit your research

Timetable

Day 1

- 10:00 - Welcome, introduction, course overview
Review of HPC skills and competencies survey
Familiarisation with fellow attendees
- 11:00 - LECTURE - Why HPC?
- 11:25 - PRACTICAL - Connecting to ARCHER
- 11:30 - BREAK – coffee/tea
- 12:00 - PRACTICAL - Sequence Alignment
- 13:00 - BREAK - Lunch
- 14:00 - LECTURE - Parallel Computing Patterns
- 14:30 - LECTURE - Measuring Parallel Performance
- 15:00 - PRACTICAL - Sequence Alignment
- 15:30 - BREAK – coffee/tea
- 16:00 - PRACTICAL - Sequence Alignment
- 16:15 - LECTURE - Building Blocks - Software
(Operating System, Processes and Threads)
- 16:45 - LECTURE - Building Blocks - Hardware
(Processors & cores, Memory, Accelerators)
- 17:15 - Review of the day
- 17:30 - Finish

Day 2

- 9:30 - Summary of day 1
- 9:45 - LECTURE - Parallel Models
- 10:30 - PRACTICAL - Fractal
- 11:00 - BREAK – coffee/tea
- 11:30 - PRACTICAL - Fractal (continued)
- 12:00 - LECTURE - HPC Architectures
- 12:30 - LECTURE - Batch Systems & Parallel Application
Launchers
- 13:00 - BREAK - Lunch
- 14:00 - PRACTICAL - Molecular Dynamics
- 15:00 - LECTURE - Compilers and Building Software
- 15:30 - BREAK – coffee/tea
- 16:00 - PRACTICAL - Molecular Dynamics
- 16:30 - LECTURE - Parallel libraries
- 17:00 - Review of the day
- 17:15 – Finish
- 19:00 – Dinner at Amber restaurant

Timetable

Day 3

9:30 - Summary of day 2

9:45 - LECTURE - Pipelines and workflows

10:15 - PRACTICAL

11:00 - LECTURE - The UK & EU HPC Landscape

11:30 - BREAK - Coffee & Tea

12:00 - LECTURE - The Future of HPC

12:30 - LECTURE - "Where next?" and things to remember

13:00 - Lunch

14:00 - Individual consultations, course review and feedback / competency survey

15:00 - Finish

Course materials

Slides, practicals, etc. available from:

<http://www.archer.ac.uk/training/course-material/2017/11/intro-epcc/index.php>