

Efficient use of the HPE Cray EX Supercomputer ARCHER2

Preliminary Agenda follows, note that a typical session includes a lecture and then practical work based on provided materials or a participants' application.

Although this event is virtual, we will aim to keep the break times free.

Tuesday 5th April 2022

9.20	Welcome and Introductions
9.30	Introduction to the Cray Hardware and Programming Environment. Focus on the HPE Cray EX hardware architecture and software stack. Tutorial on the Cray module environment and compiler wrapper scripts.
10.30	<i>Break</i>
11.00	First steps to running on Cray Hardware Examples of using the Slurm Batch system, launching jobs on the front end and first controls for job placement. Site specific environment (partitions, filesystem, reservations etc.)
~11.35	Exercises
12.00	<i>Lunch</i>
13.30	Overview of compilers and libraries An introduction to the compiler suites available on the Cray XC40. Including examples of how to get additional information about the compilation process. Special attention is given the Cray Compilation Environment (CCE)
~14.38	Exercises
15.00	<i>Break</i>
15.30	Advanced Job Launching More detailed explanation of Slurm binding technology, including advanced job placement, Multiple Application Multiple Data mode (MPMD) and thread binding.
~16.00	Exercises
17.00	<i>Close</i>

Wednesday 6th April 2022

9.00	Introduction to Perftools Overview of the Cray Performance and Analysis Toolkit for profiling applications. Focus on command line tools for low overhead performance analysis.
~9.35	Exercises
10.30	<i>Break</i>
11.00	Advanced performance analysis Part 1 Automatic performance analysis and loop work estimates with perftools. Compiler feedback and Reveal.
~11.30	Demo
12.00	<i>Lunch</i>
13.30	Advanced performance analysis Part 2 Variable scoping with Reveal. Visualization of performance data with Apprentice2.
~14.10	Performance Optimization: Improving single-core efficiency
~14.30	Demo & Exercises
15.00	<i>Break</i>
15.30	Debugging at scale Low overhead diagnostics with abnormal termination processing and stack trace analysis tool. Using gdb4hpc to launch or attach to a parallel application.
~16.00	Exercises
17.00	<i>Close</i>

Thursday 7th April 2022

9.00	Understanding Cray MPI Insight into the protocols used by Cray's MPI library and guide to optimizing communication. How to modify default behaviour and using environment variables
~09.50	<i>Mini-Break</i>
~10.00	Rank reordering and MPMD application launch
10.35	<i>Break</i>
11.00	Exercises
12.00	<i>Lunch</i>
13.30	IO Optimization -- Parallel I/O Introduction into the structure of the Lustre Parallel file system. Tips for optimising parallel bandwidth for a variety of parallel I/O schemes. Examples of using MPI-IO to improve overall application performance. Advanced Parallel I/O considerations Further considerations of parallel I/O and other APIs. Being nice to Lustre Consideration of how to avoid certain situations in I/O usage that don't specifically relate to data movement.
~14.40	Exercises or own-code porting/optimisation
15.00	<i>Break</i>
15.30	Practical Session + Open Questions and Answers Continue with Exercises or own-code optimization
17.00	<i>Close</i>