From Source to Speed: Evaluating Compilation Choices on Heterogeneous Cluster

Lead Supervisor: Dr Jianping meng

Co-supervisor: Dr Manhui Wang

Department: Research IT, IT Services, University of Liverpool

Email: Jianping.meng@liverpool.ac.uk

Project Context and Overarching Goals

Barkla2 is the University of Liverpool's recently upgraded high-performance computing (HPC) system. Like many other clusters, it delivers enhanced compute power through a heterogeneous hardware environment that includes multiple CPU architectures and GPU-equipped nodes. Such diversity makes it essential to understand how software should be compiled and managed to achieve the best performance and compatibility.

This project will examine how software compilation choices influence performance on such heterogeneous systems. Compilation translates human-written code into machine instructions, and different compilers can produce varying results depending on the target architecture. Tools such as GCC, AMD's AOCC, Intel oneAPI, and the NVIDIA HPC SDK each could have distinct behaviours. Likewise, programming languages such as Fortran and C/C++ may respond differently to these compilers and their tuning options.

Four main areas will be investigated: how different CPU compilation targets influence performance; how compiler families such as GCC, AMD, Intel and NVIDIA compare in terms of speed and compatibility; how programming languages like Fortran and C/C++ respond to varying compilation strategies; and how performance differs between software installed from pre-built packages and software compiled directly from source.

By systematically benchmarking these factors, the project will provide clear evidence on how compilation decisions shape real-world performance on heterogeneous clusters like Barkla2 as well as others. The outcomes will inform practical guidance for building and managing applications on HPC, helping researchers run their software more efficiently.

Computational Aspects and the Role of HPC

High-performance computing is central to this project, as the goal is to evaluate how software performs on a cluster with multiple CPU architectures. Our work will primarily use Barkla2, where compilation tasks will be carried out on the visualisation nodes, and performance tests will run on compute nodes in short, controlled benchmarking sessions. The results will offer practical insights into application management and optimisation only on Barkla2 but also on many similar clusters.

Interns will gain hands-on experience with HPC environments, compilers, and performance benchmarking — skills valued across academia and industry — while contributing knowledge that benefits the wider research computing community operating comparable heterogeneous clusters.