Abstract

This workshop will introduced to parallel programming for shared memory, distributed and hybrid (CPU-GPU) memory systems. OpenMP (which stands for Open Multi-Processing) is an API which supports multi-platform shared memory multiprocessing programming in C, C++, and Fortran. The Open MPI, similarly, enables the user to parallize their programs but in a different fashion. Open MPI is an implementation of MPI, which stands for "message-passing interface." MPI is "a message-passing library interface specification." More information can be found in the "MPI: A Message-Passing Interface Standard" documentation.

OpenACC is a new accelerator programming interface that provides a set of OpenMP-like loop directives for the programming of accelerators in an implicit and portable way. It allows the programmer to express the offloading of data and computations to accelerators, such that the porting process for legacy CPU-based applications can be significantly simplified. In contrast to current mainstream GPU programming, such as CUDA and OpenCL, where more explicit compute and data management is necessary, porting of legacy CPU-based applications with OpenACC requires only code annotations without any significant structural changes in the original code, which allows considerable simplification and productivity improvement when hybridizing existing applications.