

Tutorial: Developing applications with Open MPI on an OSCAR-based cluster

Jeffrey M. Squyres and Andrew Lumsdaine*
Indiana University, Bloomington, IN

Thomas Naughton and Stephen L. Scott†
Oak Ridge National Laboratory, Oak Ridge, TN

Sunday, May 15, 2005

Abstract

Although the MPI-2 specification was finished in 1996, important features of the specification and run-time environments have not become mature in MPI implementations until recently. This tutorial presents an overview of Open MPI's capabilities – and how they are useful in HPC applications – in terms of multithreading, dynamic processes, heterogeneous networking, and run-time tuning of MPI applications.

Concurrent, multi-threaded MPI applications (beyond traditional OpenMP+MPI models) with multiple application threads simultaneously executing MPI functions can be exploited for useful control and computational features. Spawning new processes and connecting to already-running MPI processes can be used for practical applications such as dynamically reporting on the status of long-running parallel codes. Using multiple networks to communicate between processes is becoming increasingly relevant, not only as organizations find that they accumulate different types of networks in LAN environments, but also in Grid / WAN environments. Finally, run-time tuning of the MPI implementation itself allows performance tweaking on both a cluster-wide and application-specific basis without changing any application code.

Emphasis will be placed on how the concepts discussed apply not only to the everyday MPI developer and user, but also to the cluster/network administrator.

Since MPI development does not require one to have systems administration skills, this tutorial will cover the installation and configuration of a HPC cluster using the OSCAR toolkit. The basic cluster concepts and administrative practices will be mentioned for those not necessarily familiar with cluster management.

The intended audience for this tutorial is the experienced MPI user who would like to learn more about the new Open MPI capabilities and would like to learn more about using OSCAR to build a HPC cluster.

*Supported by a grant from the Lilly Endowment.

†This work was supported by the U.S. Department of Energy, under Contract DE-AC05-00OR22725.