

MPI usage at NERSC: Present and Future *

Alice Koniges
National Energy Research
Scientific Computing Center
Lawrence Berkeley National
Laboratory
Berkeley, CA USA
aekoniges@lbl.gov

Brandon Cook
National Energy Research
Scientific Computing Center
Lawrence Berkeley National
Laboratory
Berkeley, CA USA
bgcook@lbl.gov

Jack Deslippe
National Energy Research
Scientific Computing Center
Lawrence Berkeley National
Laboratory
Berkeley, CA USA
jrdeslippe@lbl.gov

Thorston Kurth
National Energy Research
Scientific Computing Center
Lawrence Berkeley National
Laboratory
Berkeley, CA USA
tkurth@lbl.gov

Hongzhang Shan
Computational Research
Division
Lawrence Berkeley National
Laboratory
Berkeley, CA USA
hshan@lbl.gov

1. ABSTRACT

In this poster, we describe how MPI is used at the National Energy Research Scientific Computing Center (NERSC). NERSC is the production high-performance computing center for the US Department of Energy, with more than 5000 users and 800 distinct projects. Through a variety of tools (e.g., User Survey, application team collaborations, etc.), we determine how MPI is used on our latest systems, with a particular focus on advanced features and how early applications intend to use MPI on NERSC's upcoming Intel Knights Landing (KNL) many-core system¹ - one of the first to be deployed. In the poster, we also compare the usage of MPI to exascale developmental programming models such as UPC++ and HPX, with an eye on what features and extensions to MPI are plausible and useful for NERSC users. We also discuss perceived shortcomings of MPI, and why certain groups use other parallel programming models on the systems. In addition to a broad survey of the NERSC HPC population, we follow the evolution of a few key application codes² that are being highly optimized for the KNL architecture using advanced OpenMP techniques. We study how these highly optimized on-node proxy apps

¹Cori Phase-2 system will be a Cray XC system based on the second generation of Intel Xeon Phi Product Family, called Knights Landing (KNL) Many Integrated Core (MIC) Architecture. The system will have a sustained performance that is at least ten times that of the NERSC-6 "Hopper" system, based on a set of characteristic benchmarks.

²<http://www.nersc.gov/users/computational-systems/cori/application-porting-and-performance/application-case-studies/>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

EuroMPI '16 September 25-28, 2016, Edinburgh, United Kingdom

© 2016 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-4234-6/16/09.

DOI: <http://dx.doi.org/10.1145/2966884.2966894>

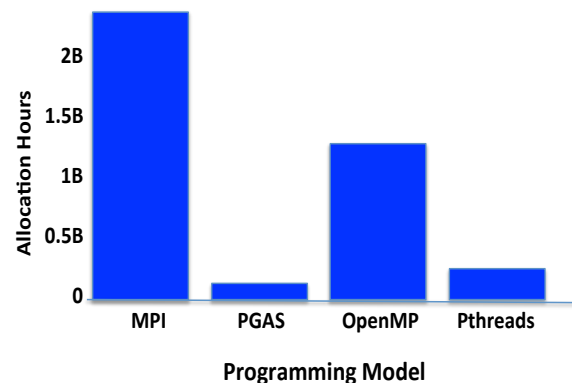


Figure 1: Approximate distribution of parallel programming models used at NERSC based on 2015 allocation requests to US Department of Energy. Total exceeds 100% because some codes use multiple models.

and full applications start to make the transition to using full hybrid MPI+OpenMP implementations on the self-hosted KNL system.

CCS Concepts

•Software and its engineering → General programming languages;

Keywords

MPI, programming models

2. ACKNOWLEDGMENTS

This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.