

Programming Models & Environments Summit

February 5, 2014

American Geophysical Union

2000 Florida Ave. NW, Washington D.C. 20009

Summit Goal:

1. Discuss the major challenge of creating a system-wide programming model that supports heterogeneous processing elements that are distributed along the data path within an extreme scale computer.
2. Create a vision for a programming environment that enables high performance and energy efficient computations, application performance portability and “ease of programming” for exascale systems and beyond.
3. Discuss the need for an extreme scale programming workshop.

Background:

Using current programming models results in the intractable, manual navigation of a complex space and in codes that are not performance portable. Writing a single program with MPI+{OpenMP, CUDA, and OpenACC} is extremely complex and time consuming, requiring understanding and coordinating between different programming models, syntaxes, and even languages. In addition, programmers today need to decide how to distribute work on to a set of accelerators and other local and global resources. These programming models and approaches have a costly code portability consequence: applications will need to be re-written in order to run on different computer systems. And when considering performance portability, the costs of tuning may be unacceptably high.

The proposed scope of research needed in programming models for Exascale and beyond starts at the high-level representation of an application and engages the automated support for optimizations and mappings of these high-level specifications to abstract machine representations and optimized dynamic runtime codes for particular platforms. These automated mappings and optimizations support performance portability of code so that it is “write once, run anywhere.” The different classes of Exascale programmers that will be part of the loop generating optimized runtime code to a specific platform need to be considered as we conceive a new multi-level programming model environment. We need to achieve a unified programming model for processing along the intensely heterogeneous data path, which includes accelerators, processing in memory and near memory, processing in storage, processing in network interfaces, etc. Legacy application codes will create a significant challenge when ported to future extreme scale computers.

The proposed unified, system-wide programming model for extreme scale computers wouldn't just be inclusive of hardware heterogeneity that is anticipated for Exascale: it needs to be flexible enough to include also future, “beyond Exascale,” hardware and software features.

Summit Questions

We need to hear your views on how to best address the challenges that we face in this area. We also need to discuss the state-of-the-art approaches that can be used to solve them. The following questions will be discussed during this meeting:

1. What are the main issues that Extreme scale programming models need to address and why current programming models do not address them?
2. What is the vision and set of goals for a unified, system-wide programming model for extreme scale computers that enables high performance and resilient, energy efficient computations, application portability and “ease of programming?”
3. What are the elements of a research program that will enable a unified, system-wide programming model to become prevalent in the 2022 time frame? What other activities will be needed in order to ensure this goal?
4. What current research efforts support the development of a unified, system-wide programming model?
5. How will “legacy” application codes be supported in this new programming model?

Summit Agenda

We plan to start at 8:30 a.m., break for lunch from 11:30 – 12:30 p.m., returning for the 12:30 p.m. – 5:30 p.m. afternoon session. Breakfast will be included with your registration and available at 7:30 a.m.

8:30 – 8:45 a.m.	Brief summit introduction (Sonia) Our concerns (Harrod, Thuc?)
8:45 – 9:35 a.m.	Initial thoughts of each participant (5 minutes per participant)
9:35 – 10:15 a.m.	Discussions to refine/redefine set of questions to be discussed/answered
10:15 – 10:30 a.m.	Break
10:30 – 11:30 a.m.	For question 1, comments from participants (30 minutes) and discussions (30 minutes)
11:30 – 12:30 p.m.	Lunch break
12:30 – 1:30 p.m.	For question 2, comments from participants (30 minutes) and discussions (30 minutes)
1:30 – 2:30 p.m.	For question 3, comments from participants (30 minutes) and discussions (30 minutes)
2:30 – 2:45 p.m.	Break
2:45 – 3:45 p.m.	For question 4, comments from participants (30 minutes) and discussions (30 minutes)
3:45 – 4:45 p.m.	For question 5, comments from participants (30 minutes) and discussions (30 minutes)
4:45 – 5:30 p.m.	Wrap-up, discussion about follow up workshop and report assignments

Summit Participants:

Research community:

Michael Garland
Mary Hall
Michael Heroux
Rich Lethin
Jim Kahle (for Kathryn O'Brien)
Vivek Sarkar
John Shalf
Marc Snir
Thomas Sterling
Kathy Yelick

DOE program managers:

William Harrod
Thuc Hoang
Sonia Sachs
Lucy Nowell
Thomas Ndousse-Fetter

Proposed Summit Report Outline

1. Executive Summary
2. Background: Challenges to be addressed
3. The top five questions to be addressed, including alternative solutions, advantages and disadvantages
4. Road Map and Conclusions