

PGAS Compilers for Community Codes Development: Are We There Yet?

Panel

Chair: Sadaf Alam, Swiss National Supercomputing Centre

Abstract:

Today community driven scientific libraries and applications for domains such as climate, chemistry, material, engineering and life sciences utilize a large fraction of resources on leadership supercomputing platforms as well as institutional clusters. Majority of these applications are implemented in Fortran, C and C++ with message-passing MPI parallel programming paradigm. Despite the standardization and potential performance and productivity gains that PGAS languages (CAF and UPC) offer, there are still obstacles to widespread adoption of these languages as scientific code developers seek code and performance portability across a range of systems. Panelists in this session represent compiler developments efforts from industry, government labs and academia (Cray, IBM, Intel, LBNL, Rice University and University of Houston). They will present their respective compilers' roadmaps and discuss what measures are needed for encouraging usage of PGAS languages among scientific communities on emerging HPC systems.

Panelists:

Gheorghe Almasi, IBM
Brad Chamberlain, Cray
Barbara Chapman, University of Houston
Ronald Green, Intel
Costin Iancu, LBNL
John Mellor-Crummey, Rice University