

Parallel Programming using the PGAS Approach

Friday, September 5, 2014 9:00 AM (40 minutes)

The two most common approaches for parallel programming are message passing (for example using MPI, the message passing interface) and threading (for example using OpenMP or Pthreads). Threading is generally considered an easier and more straightforward solution for parallel programming but it can generally only be used on a single shared memory node. MPI, on the other hand, scales to the full size of today's machines, but it requires a more complex planning and orchestration of data distribution and movement.

PGAS (Partitioned Global Address Space) approaches try to combine the best of both worlds, providing a threading abstraction for programming large distributed memory machines. Data locality is made explicit in order to be able to take advantage of it for performance and energy efficiency reasons. The talk will give an introduction to the concept of PGAS programming and provide examples using UPC (unified parallel C). The research project DASH, which provides a realization of the PGAS model in the form of a C++ template library, will also be introduced in the talk.

Presenter: FÜRLINGER (UNIVERSITY OF MUNICH), Karl

Session Classification: Plenary talks