Parallel acceleration: achievements and pitfalls

Gundolf Haase[1]

Medium size computer clusters with several hundred CPU-cores are already available to small and medium size companies. The bunch of cores doesn't perform always as well as expected for several reasons and these miracles continue when GPUs are used as additional performance booster.

Based on several examples we will show great parallel (strong) speedup on multiple cores and on clusters of compute nodes as well as (a few) discouraging results that might have to be solved by means of mathematics. Especially the GPU parallelization combines all sorts of pitfalls from shared and distributed memory computing.

We will present several student projects from physics and engineering and how they advanced during the winter term. The examples range from non-Newtonian fluids for shared memory systems, or PDE solvers on GPU to a (random walk) worm algorithm for Lattice-Boltzmann on GPU. The performance gain on GPUs ranged from dissatisfying (factor 3) for a standard package to splendid (factor 300) in case of quantum chromodynamic problem. The typical acceleration by a GPU is a factor of 10 for memory intense algorithms and a factor 50 and more for compute intense algorithms.

References:

[1] Ph.D. Seminar on High Performance Computing, Graz, Winter 2010/11, http://www.uni-graz.at/~haasegu/Lectures/GPU_CUDA/WS10/projects.html

[1] University of Graz, Institute for Mathematics and Scientific Computing, Graz, Austria, gundolf.haase@uni-graz.at