Scientific Computing & Data Analysis Section

Introduction to CUDA and GPU Programming

Graphics cards — GPUs — are a common way to greatly increase the speed of heavy computations, both on clusters and on personal workstations. If you have the right kind of workload, you can improve your speed by five or ten times, and sometimes even more.

In this course we give a hands-on introduction to CUDA, a C language extension for programming NVIDIA GPUs, and show how to build and run GPU code on the OIST clusters. You need no experience with parallel programming, but you should have some brief previous experience with C or C++.

Next Session

Monday, March 11, 2019, 13:00-15:00 in room C209

Course material

Presentation slides: CUDA presentation [PDF]
Slides and Example code: GPU.zip