CUDA for scientific computing

10th and 11th of September in Espoo

The CUDA for scientific computing course consisted of lectures by Johan Hedborg (Synective Labs Ltd.) and hands on exercises. The course made the gains and limitations of performing calculations using the *general purpose graphics processing unit* (GPGPU) rather clear. The hardware used for the practical work was brand new Nvidia Tesla graphics cards.

The gains were made obvious by an example where moving simple non-optimized code from the top-of-the-line *central prosessing unit* (CPU) to GPGPU environment provided roughly 5-times faster calculations and optimized code 12-times faster calculations. The downside is that the programming language is a low-level language, which makes programming more laborious than with some higher level languages available for CPUs.

CUDA (and OpenCL) make it possible to do massively parallel calculations on a desktop computer very quickly. Although, the actual supercomputers are still far superior (Theoretical peak performance for CSC Louhi cluster is ~102 Tflops/s; one quad-tesla desktop ~3 Tflops/s) the Tflops per euro or per kWh ratio is much better on GPGPU based clusters.

Jani Seitsonen

Institute of Biotechnology, University of Helsinki