

Programming on the GPU with CUDA

Two-day course for researchers, PhD and MSc students

November 11-12, 2024

What is a GPU? Nowadays the Graphics Processing Unit (GPU) is a mainstream hardware component in high-performance computing. For affordable budgets anyone can harness supercomputer performance. However, realizing efficient parallelism combines three knowledge areas: firstly, on the architecture and compute capabilities of the GPUs; then, on special constructs for programming a GPU-equipped computer; finally, on the special algorithms for performing logical and mathematical operations in parallel.

What is CUDA? CUDA stands for Compute Unified Device Architecture. It is a software-development tool kit for programming on the GPUs maintained by the mainstream manufacturer Nvidia. CUDA provides language extensions for C, C++, FORTRAN, and Python as well as knowledge-specific libraries. A single source code is then able to instruct the CPU and GPU alike. Also, CUDA-extended codes keep pace closely with the rapid developments in the underlying technology.

Goals and prerequisites To guide you in this development niche, the [Delft Institute for Computational Science and Engineering](https://www.tudelft.nl/en/department/computational-science-and-engineering/) (DCSE) offers a CUDA course every quarter. We will explain basic principles and advanced topics on GPU programming with CUDA. You will apply these notions in our labroom with hands-on examples. After this course you will be able to get simple CUDA programs running on a GPU-equipped computer. As prerequisite, a rudimentary understanding of programming languages like C++ or Java is ideal; that of FORTRAN or Python will be helpful too. Some interest in linear algebra and iterative solvers is a little advantage.

Schedule and instructors Prof Kees Vuik, Ir Kees Lemmens, Dr Matthias Möller of the Faculty of Electrical Engineering, Mathematics and Computer Science teach the course. For the schedule we refer to <https://www.tudelft.nl/evenementen/2024/dcse/courses/programming-on-the-gpu-with-cuda>.

Location Delft University of Technology, Faculty of Mathematics & Computer Science, Mekelweg 4, 2628 CD Delft. The classes are given in Building 36, 'Penguin' Laboratory Room (HB02.130).

Costs and registration For members of DCSE this course is free. TU Delft staff and students pay € 50 for a single day and € 100 for 2 days, for other participants the costs are € 200 for 1 day or € 350 including lunch and course material. Lunch, refreshments, lecture materials are included.

Please sign up at <https://www.tudelft.nl/evenementen/2024/dcse/courses/programming-on-the-gpu-with-cuda>

The maximum number of participants is 25.

More info The course's splash page is <https://www.tudelft.nl/evenementen/2024/dcse/courses/programming-on-the-gpu-with-cuda>

Please contact C.W.J.Lemmens@tudelft.nl for content or dcse@tudelft.nl for administration and logistics.

