

Introduction to Programming on the GPU with CUDA

organized for

Researchers and Students (PhD, MSc)

April 12-13, 2018

The Graphics Processing Unit or GPU is nowadays a mainstream component in Scientific Computing. For relatively little money one can have supercomputer performance. However, some extra work has to be done to make an ordinary sequential program suitable for use on the GPU.

One of the most important tools for using GPUs is currently “CUDA” (Compute Unified Device Architecture). This is basically an extension to the C programming language, which can be used to program the GPU in an easy way. Another tool is OpenCL, but as this is less mature than CUDA and performs less on especially the newest GPUs, the focus will be currently on CUDA.

Course goal In this 2-day course we will explain the basic principles of GPU programming and we will let you practise with many examples on several types of GPUs in our labroom. After this course you should be able to make simple CUDA programs which can be run on a GPU.

Preliminaries We recommend having a (rudimentary) understanding of a C-like programming language, such as C++, Java or similar. Fortran or Python should be fine as well. Interest in iterative solvers can be helpful.

Teachers Prof.dr.ir. Kees Vuik, Ir. Kees Lemmens, and Dr. Matthias Möller

Costs For members of DCSE this course is free, TU Delft staff and students pay € 50,- for a single day and € 100,- for the full 2 days course. For other participants the costs are resp. € 200,- for 1 day and € 350,- for 2 days. This is mainly to cover expenses for the lunch and the course material.

Location EEMCS building, Penguin Lab, van Mourik Broekmanweg 6, Delft.

Schedule

	Day 1 Introduction		Day 2 Advanced	
09:15–09:30	Arrival, coffee, tea		Arrival, coffee, tea	
09:30–10:30	Parallel computing on GPUs	KV	Solvers	KV
10:45–11:30	GPUs: design and architecture	KL	Lab 4: Shared memory, streams, atomics	KL
11:45–12:30	Lab 1: CUDA introduction	KL	Lab 5: Optimising code: Innerproduct	KL
12:45–13:30	Lunch		Lunch	
13:30–14:45	Lab 2: Using CuBlas, CuFFT	KL	Lab 6: Unified memory, Dynamic Parallelism	KL
15:00–16:30	Lab 3: Debugging and Profiling	KL	Applications on GPU	MM
16:30–17:00	Lab 3: Debugging and Profiling	KL	Lab 7: Applications on GPU	MM

Notes The number of participants is restricted to 20. This course is given roughly every quarter.

Register In order to attend this course please register at <https://www.aanmelder.nl/gpucourse2018>.

More info <http://www.cse.tudelft.nl>.

Email contact: C.W.J.Lemmens@tudelft.nl or D.M.Dongor@tudelft.nl

