

Fundamentals of Accelerated Computing with CUDA C/C++

February 5, 2020

VSB - Technical University of Ostrava

IT4Innovations

Lecturer: Dr. Momme ALLALEN | Leibniz Supercomputing Centre, Germany

The CUDA computing platform enables the acceleration of CPU-only applications to run on the world's fastest massively parallel GPUs. You experience C/C++ application acceleration by:

- Accelerating CPU-only applications to run their latent parallelism on GPUs
- Utilizing essential CUDA memory management techniques to optimize accelerated applications
- Exposing accelerated application potential for concurrency and exploiting it with CUDA streams
- Leveraging command line and visual profiling to guide and check your work.

The lectures are interleaved with many hands-on sessions using Jupyter Notebooks. The exercises will be done on a fully configured GPU-accelerated workstation in the cloud.

This training is a part of **NVIDIA AI & HPC ACADEMY 2020**.

AGENDA

- 09:00-09:20 | Introduction
- 09:20-11:00 | Part1: Accelerating Applications with CUDA C/C++
- 11:00-11:15 | Coffee Break
- 11:15-13:00 | Part1: Continued
- 13:00-14:00 | Lunch
- 14:00-15:30 | Part2: Managing Accelerated Application Memory with CUDA Unified Memory and nvprof
- 15:30-15:45 | Coffee Break
- 15:45-16:50 | Part3: Asynchronous Streaming and Visual Profiling for Accelerated Applications with CUDA C/C++
- 16:50-17:00 | Q&A, Final Remarks



More information & registration:
events.it4i.cz/event/40/



This event was partially supported by The Ministry of Education, Youth and Sports from the Large Infrastructures for Research, Experimental Development and Innovations project "e-Infrastruktura CZ – LM2018140" and partially by the PRACE-6IP project - the European Union's Horizon 2020 research and innovation programme under grant agreement No. 823767.