

Title:	Parallel and Distributed Computations
Lecture hours:	30
Study period: (summer/winter)	summer
Number of credits:	3
Assessment methods:	Tests and reports
Language of instruction:	English
Prerequisites:	English B1/B2, basics of programming, knowledge of C/C++, computer architecture
Course content:	<p>As since March 2011 (the first in Poland) until April 2017 UKW was NVIDIA GPU Education Center, and as still the teaching activity of the university is supported by the NVIDIA, the content of the course is partially based on the NVIDIA company teaching materials.</p> <p>The subject of the lecture are selected problems of parallel and distributed computing with special emphasis on the GPU and CUDA C computing. During the lecture the following details about CUDA C are explained: principal concepts of CUDA C (what is host, device, kernel, block, grid, thread, warp); function declarations; launching program in CUDA C; memory models and types in CUDA (registers, shared memory, global memory, constant memory); data transfer (host to host, host to device, device to host, device to device); performance analysis of CUDA programs.</p> <p>Students perform practical tasks for CUDA C programming. The tasks are following: setting up IDE environment for CUDA C programming; launching the first program in CUDA („Hello world”); testing principals of CUDA language, launching kernel function; launching code for testing and explaining the difference between CUDA kernel, host and device function (keywords: <code>_global_</code>, <code>_device_</code>, <code>_host_</code>); launching code for testing data transfer and memory allocation (<code>cudaMalloc()</code>, <code>cudaMemcpy()</code>, <code>cudaFree()</code>); code for checking the GPU device properties with the the use of <code>cudaGetDeviceCount</code> and <code>cuda-GetDeviceProperties</code>; launching the code for vector adding with the use of blocks for parallel computing;</p> <p>launching the code for vector adding with the use of threads; launching the code for scalar product of two vectors with the use of shared memory (<code>_shared_</code>).</p>
Learning outcomes:	<p>Students will learn how to program heterogeneous parallel computing systems and achieve high performance of the computations.</p> <p>Students will understand:</p> <ul style="list-style-type: none"> – how to apply parallel programming API, tools and techniques principles of parallel algorithms – GPU architecture features and constraints <p>Students will gain comprehensive knowledge in contemporary issues of the parallel and distributed computations.</p>
Name of lecturer:	PhD Eng. Sebastian Kula
Contact (email address):	skula@ukw.edu.pl

Literature:	<ol style="list-style-type: none">1. Programming massively parallel processors: a hands-on approach, David B. Kirk, Wen-mei W. Hwu., Morgan Kaufmann Publ., 2010.,2. CUDA by example: an introduction to general purpose GPU programming, Jason Sanders, Edward Kandrot, Addison-Wesley, cop. 2011.3.NVIDIA CUDA ZONE https://developer.nvidia.com/accelerated-computingtraining
--------------------	---