

GPU Programming 2017/18

Tutorial 7

In this tutorial we will write a simple Cuda program that adds a constant to a vector `data`.

- 0.) Download the [skeleton code](#) and generate the build system using `cmake`.
- 1.) Implement the steps required on the host to run a Cuda program:
 - i.) Query the number of Cuda devices, check the required properties, and initialize a suitable one.
 - ii.) Allocate memory on the device and copy the input data to the device.
 - iii.) Set up the number of thread blocks and threads that are required and call the device program.
 - iv.) Implement the kernel program.
 - v.) Copy the result back to the host and clean up device memory.

Please finish the implementation until next week (week of 6/11/2016).