GPU Programming 2019/20

Tutorial 3

In this tutorial we will consider the question of correctness of parallel programs for floating point operations. As example problem we will again consider the computation of the norm of a vector $a \in \mathbb{R}^n$.

0.) Download the skeleton code and generate the build system using `cmake`.

1.) Allocate the vector a of type `DataType` and initialize it to

$$a_i = \frac{1}{i+1}. \tag{1}$$

Implement the function `computeNorm()` in the skeleton code using the pool pattern that has been discussed in class.

2.) Compute the norm of a with 1, 2, and 4 threads. How can you reasons about the correctness of your implementation even when different interleavings yield different numerical results?

Please finish the implementation until next week (week of 13/11/2016).