Parallel Programming 2013/14 - Lab exercise 2

Building on the previous exercise (simple array reversal), this exercise introduces multi-block kernels and shared memory. The content has been adapted from that of the nVidia CUDA training website [1].

1. Download the program skeleton:

   `http://www.cs.bham.ac.uk/~hxt/2013/parallel-programming/lab2a.cu`

   Your task is to complete the kernel code. You may want to draw a diagram to understand the distribution of work among multiple thread blocks.

2. Download the program skeleton:

   `http://www.cs.bham.ac.uk/~hxt/2013/parallel-programming/lab2b.cu`

   Your objective is to fill in the missing code (follow the comments). You may find the following diagram useful:

   `http://www.cs.bham.ac.uk/~hxt/2013/parallel-programming/shared_mem.pdf`

3. Without using any profiling tools, can you suggest (and implement) a method to compare the performance of the above two programs?

References

[1] nVidia. CUDA Zone - Education and Training. Available at `https://developer.nvidia.com/cuda-training`, 2013.