

CS5234 Advanced Parallel Computation Spring 2013

Homework 1: Introduction to CUDA

1 Due Date

Homework 1 is due on Wednesday February 6th, 2013, 11:59pm.

2 Introduction

This homework is designed to help you understand CUDA programming framework and practice on your first CUDA program. You need to learn the CUDA SDK template project, how to compile and debug a CUDA project.

3 Project requirement

Please write a CUDA program compute a dot product between two vectors. Please implement the multiplication of every two elements from the vectors on GPU using CUDA. The final summation should be calculated on CPU.

A dot product between two vectors $\vec{a} = [a_1 a_2 \dots a_n]$ and $\vec{b} = [b_1 b_2 \dots b_n]$ with length n can be definite as

$$c_i = a_i \times b_i \quad (1)$$

$$c = \sum_{i=1}^n c_i. \quad (2)$$

The first equation/step is computed on GPU using CUDA and the second equation and step is computed on CPU.

Please use a random number generator to create the elements in each vector (For example, using `rand()` function.) Each element is a float point number, ranging between 0.0 and 1.0.

Please also use three different lengths for the vectors, $n = 1024, 65536, 1048560$. Calculate the dot product results for all three vectors.

When you configure the number of threads for each block, please enumerate different numbers, from 32 to 512, incrementing 32 each time.

Please record the execution time for the CUDA kernel function call, the data transfer time (including both copy to GPU and from GPU), and the overall calculation time. Record these timing information for each input and each thread configuration. Be careful about how to record the timing information using CUDA timer, please refer to the section 5.1 (Timing) of “*CUDA C BEST PRACTICES GUIDE*” or the section 3.2.5.6(Events) of “*NVIDIA Cuda C Programming Guide Version 5.0*”.

Please also implement a pure CPU version of the dot product. In each execution of your GPU version, please also compare with the CPU version to make sure the final result is the same.

Please use Excel to draw the figures of the timing information, one figure for each input. The X axis is the number of threads per block, the Y axis is the CUDA kernel execution time.

4 What to Submit

Put your solution in one or more source files. The main file (which includes function `main()`) should be named `homework1.cpp` or `homework1.cu`. Include all your source files in a zip file and upload to class Scholar site in your own dropbox. Please do not forget to include your excel figure files (in JPG format). Please do NOT include the compiled EXE files.

Note: Please use CUDA 5.0 for all the homeworks for this course.