CS6023 GPU Programming
Problem Set 1: Computation

1. Write a CUDA program to add two vectors of the same size. Modify the program to make it work even if the vector sizes are different.
2. Write a program to find out the unique id of each thread irrespective of what kernel launch configuration is used (that is, single / multiple blocks, number of dimensions of threads / blocks, etc.).
3. Create an array of strings (character arrays or STL strings) on the CPU. Transfer these to the GPU and print the strings from various threads.
4. Write a kernel with if condition but without any thread-divergence.
5. Write a kernel that takes at least one parameter d where d is the degree of divergence in the range 0..32. 0 indicates no divergence, while 32 indicates that all the warp-threads execute different instructions. Plot the times taken by this kernel for various values of divergence.