

# Homework 4

## Required:

1. Do you find any discrepancy between what I taught in the “Manycore Computing” lecture and what <https://cvw.cac.cornell.edu/gpu/tesla> says about Nvidia Tesla K20 GPU? Which one do you think is correct? Why?

## Bonus:

2. Write a device query program to query the capabilities of a CUDA-capable GPU. Run it on a GPU node on Hyades and report the results.
3. Write a CUDA program to perform Matrix-Vector multiplication.