

Due Feb 17

1. Get CUDA to work for you. (depending on your architecture, mode of access etc. this means different things to different people)
2. Run a few of the demo examples in the CUDA SDK (if you have non-graphics access you may not be able to run those applications that require graphics). Report on the outputs of a few of the applications.
3. Write the simple vector and matrix addition skeleton examples in Chapter 2 of the programming guide as actual functions.
From the programming guide figure out a way to time your function and provide a table of time versus matrix size. Also, implement a CPU version of your function, and compare the results with the GPU output. Note your system/CPU/GPU version and anything else that is relevant

Due Feb 24

4. Explore matrix matrix product routines on the CPU. (Don't spend too much time on Strassen's algorithm). Write and time a simple direct matrix-matrix products on the CPU. Be aware of cache and loop ordering effects
5. Look at the matrix-matrix product example in Chapter 6 and get it to work. This will be discussed next week by one of you in class. Time it, and compare it with the GPU version.
6. Compare this with the CUBLAS version.