Senior Honors Seminar (GPGPU programming) Spring, 2008
CS U610 Professor G. Cooperman
http://www.ccs.neu.edu/course/csu610/

Course Description: Please note that this course is cross-listed with ECE U630, for which Prof. D. Kaeli is the instructor of record. In fact, the course will be taught jointly.

This course will explore what are the current limits of high performance computing. The limits include CPU power, RAM speed, size of RAM and network bandwidth. However, recently, one of those limits (for commodity computing) had its ceiling drastically raised. The general term used is GPGPU (General Purpose GPU programming, where GPU is Graphics Processing Unit).

GPGPU is posed to create a revolution in high performance computing. The latest evolution of the graphics processor of a video board makes it closer than ever to a general purpose CPU. Note especially the G8 series (GeForce 8800) from nVidia. The GeForce 8800 has 128 CPUs, spread over 8 GPU chips, that require a style of programming called SIMD (which was popularized by Intel in their SIMD instructions for the Pentium). For example, in regular computatins such as matrix multiplication, times 100 times faster than what the CPU is capable of have been reported. ATI is expected to soon announce a competing model.

This course will quickly describe the programmer's model for the G8 processor. (However, the model is more subtle than a completely general CPU.) We will gain a first experience programming is using the CUDA simulator. We will then test the practical limits by running on the actual hardware.

We will then test various computational kernels (the inner part of a computation that requires 90% of the CPU time). By default, we will be organized in teams of two, although other configurations are possible. Because GPGPU is so new, not all computational kernels have yet been tested. Hence, there is an opportunity to take some of our results, and spend additional effort (probably after the Spring) to turn it into a paper at a respected conference.

In evaluating GPGPUs, it is important to compare it with competitive approaches. In addition, GPGPU programming shares some features with these competitive approaches, and may even give rise to hybrid models, such as use of multiple GPGPUs. For this reason, we will also have a parallel track in which the professors will lecture on pertinent background in parallel computing.

Faculty Information:

Professor G. Cooperman
Office: 336 West Village H
e-mail: gene@ccs.neu.edu
Phone: (617) 373-8686
Office Hours: Tues. and Fri. at 5:40 – 6:40, and by appointment.

Textbooks:

None. (We will use materials from Web.)

Exams and Grades:

The course will be taught in a seminar style. There will be a course project, done by teams of two. Each team will give an oral report, and write an interim written report and a final report.

The TAs will be available to help with difficulties in the programming. The reports will demonstrate whether students have absorved the lessons provided by the TAs and can clearly describe the driving principles and techniques of their project.

The course grade will be determined 25% by the oral report, 30% by the interim report, and 45% by the written report. Class participation will provide up to an additional 10% of "extra credit".

As a seminar course, it would be my hope that everyone will receive an A, although as always, that will depend on the individual performance of the students.