CS 5984
Advance Computer Graphics:
General Purpose Computing and Visualization on GPU

Introduction

In recent years, the peak performance of a GPU has exceeded that of the CPU by several orders of magnitude. Modern GPUs have a lot more computational cores than top-of-the-line CPUs, and show great performance increase for many data parallel applications, such as image/video processing, scientific simulation and visualization. The commoditized GPU also exhibits advantages over other computing devices on cost-performance ratio. Therefore, many commercial applications and research projects have leveraged the computational power of GPU for performance enhancement. In this course, we will learn a GPU programming model, NVIDIA's CUDA, and the basic programming concepts for general purpose computing and visualization on NVIDIA GPUs.

Pre-requisite

- CS4204 Computer Graphics
- CS4234 Parallel Computation

Course Objectives:

- Understand the system architecture, especially memory hierarchy, of GPUs.
- Understand programming model and language specifications of GPU programming.
- Understand how to transform serial code that is amenable to data parallelism into GPU-accelerated code.
- Learn the optimization strategies for GPU-base programming.
- Learn the basic visualization techniques, such as OpenGL-based rendering and ray-tracing, on GPUs.
- Learn the basic data sharing techniques between a CUDA program and a OpenGL-based visualization program.

Course Work

Below is an *estimate* of the contributions of different parts of your final grade.

- Participation & Quizzes 10%
- 3-4 Homeworks 30%
- Midterm Project 25%
- Final Project & Report 35%

Midterm project is about the CUDA implementation of a general purpose computing application, such as data mining or image processing.

Final project is about a scientific simulation application together with GPU-based visualization of the simulation results.

Materials and References

There is no required Text Book for this course. We use the literatures from conference and journal papers. You can use the following list as reference materials:

1. NVIDIA CUDA 2.x Programming Guide, 2009.
2. NVIDA CUDA website, http://www.nvidia.com/object/cuda_home.html.
3. UIUC Parallel Programming Course Website: http://courses.ece.uiuc.edu/ece498/al/.
4. Astro GPU Workshop Videos: http://www.astrogpu.org/videos.php
5. SC07 GPU Tutorials on GPGPU.org Website: http://www.gpgpu.org/sc2007/

6. Other Courses on GPGPU at NIVIDA Website: http://www.nvidia.com/object/cuda_university_courses.html
7. GPU Computing Course at SIGGRAPH Asia: http://sa08.idav.ucdavis.edu/

List of Topics

- GPGPU Computing
- Shaders and General purpose computing
- CUDA Programming Model
- Computation to Core Mapping
- GPU Memory (Shared Memory Example)
- GPU Memory II (Memory Hardware and Bank Conflict)
- Optimization Strategies (Global Memory Access Pattern and Control Flow)
- Multi-GPU computing
- OpenGL inter-ops with Pixel Buffer Object (PBO) and Vertex Buffer Object (VBO)
- Ray-tracer on GPUs