

## Performance Optimization of Applications on GPUs

**Professor:** J. (Ram) Ramanujam, 345 EE Bldg., 578-5628 Email: jxr@lsu.edu  
**Time:** 5:10pm - 6:30pm MW, 149 EE  
**Reference:** *Class notes and recent papers from the literature*  
**Prerequisites:** Graduate standing

### Description (EE 7700-1)

The course deals with the study of optimizations aimed at improving performance of science and engineering applications on GPUs and heterogeneous systems consisting of CPUs and GPUs. Basic ideas of parallel computing will be covered first. The course will then discuss the two GPU programming models, CUDA and OpenCL. Next, we will discuss higher-level programming approaches for GPUs such as PyCUDA, HiCUDA, HMPP and Accelerator. In addition, we will discuss optimization strategies, application performance optimization and the use of compiler tools such as PLUTO for automatic code generation. Program transformations that enhance data locality and optimize memory usage will get a lot of coverage. In addition, compilation starting from high-level specifications or for domain-specific languages will be discussed.

### Course Outline

#### 1. Introduction

- Parallel computing
- Multicore processor architectures, GPGPUs
- CPUs versus GPUs
- Brief Overview of Programming Models for CPUs and GPUs
- Brief Overview of Optimizations for CPUs and GPUs

#### 2. GPU Architectures and programming

- GPU architectures
- General-purpose computation on GPUs
- Programming models and idioms
- GPU programming models/environments:
  - OpenCL
  - CUDA
  - PyCUDA
  - HiCUDA
  - PGI Accelerator
  - CAPS/HMPP
- Code examples on GPUs
- Examples of CPU vs. GPU performance

#### 3. Optimizations and tuning for GPUs

- Performance characterization
- Performance profiling
- Application case studies
- Optimizing memory accesses
- Multi-level parallelism exploitation
- Performance models and empirical search
- Compiler-driven tuning
- Examples of application optimization
- Software managed memory hierarchies
- Optimizations in HMPP, Accelerator, PLUTO

#### 4. Optimizations and tuning for heterogeneous CPU-GPU systems

- Performance characterization
- Application case studies
- Examples of application optimization