

CS 706: PARALLEL PROGRAMMING ON HETEROGENEOUS PLATFORMS (3-0-0:3)

Basics of Concurrency Theory

Race Conditions, Mutual Exclusion with Semaphores.

Scheduling Algorithms

Multicore Architecture, SIMD execution model.

GPUs as Co-Processors for Heterogeneous Computing.

Nvidia GPU architecture (Fermi and Kepler Architecture)

CUDA programming model. Introduction to Heterogeneous Programming with CUDA.

Memory model of GPU. Tiling for conserving memory bandwidth.

Introduction to OpenMP, OpenCL and C++AMP.

Exercise programs in CUDA, OpenMP, C++AMP, OpenCL.

References:

1. A.W. Roscoe, The Theory and Practice of Concurrency Theory.
2. M.J. Flynn, Computer Architecture: Pipelined and Parallel Processor Design, Narosa Publishing House/ Jones and Bartlett.
3. D.A. Patterson, J.L. Hennessy, Computer Architecture: A Quantitative Approach, Morgan Kaufmann Publishers.
4. David Kirk and Wen-mei Hwu, Programming Massively Parallel Processors: A Hands-on Approach (Applications of GPU Computing Series).
5. <http://docs.nvidia.com/cuda/cuda-c-programming-guide/index.html>