Introduction to Parallel Computing

平行計算概論

Instructor: **Prof. Chun-Yi Lee, Spring**
Location: Delta Building R104

Course Outline and Goal

From smart phones, to multi-core CPUs and GPUs, to the world's largest supercomputers, parallel processing has become the mainstream high-performance processor architecture. The goal of this course it to provide a deep understanding of the fundamental principles and engineering trade-offs involved in designing modern parallel computing systems. This course is offered for **junior** and **senior undergraduate students**. The following topics will be covered in the course:

Part I: Parallel programming languages (6 weeks)
- Multi-thread programming
- Message passing interface (MPI) programming
- OpenMP programming
- GPU architectures and CUDA programming

Part II: Parallel programming algorithms and concepts (5 weeks)
- Parallel execution and efficient execution
- Task decomposition and load balancing
- Task scheduling and distribution
- Synchronization models
- Performance evaluation: performance modeling and simulations

Part III: Parallel computer system and architectures (5 weeks)
- Parallelism at the processor level
- Shared memory and distributed memory
- Heterogeneous systems
- Cache coherence
- Memory consistency
- On-chip interconnection networks (NoCs)

Lecture Hours

Tue: 3:30pm - 5:20pm (two hours)
Thu: 3:30pm - 4:20pm (one hour)

Prerequisites

- Introduction to programming (I) (II)
- Data structures
- Digital logic design
- Computer architecture

Assignments

- Two MPI programming assignments
- Two CUDA programming assignments
- One cache coherence modeling and simulation assignment
- One parallel system performance simulation assignment

Textbook

There is no textbook for this course. Handouts will be provided online for students to download. The following books are suggested as reference materials.

- **Principles and Practices of Interconnection Networks**

  by William James Dally, Brian Patrick Towles
- **Parallel Computer Architecture: A Hardware/Software Approach**

  by David Culler and Jaswinder Pal Singh
- **Memory Systems: Cache, DRAM, Disk**

  by Bruce Jacob, Spencer Ng, and David Wang

Grading

- Assignments (6 times): **50%**
- Midterm exam: **20%**
- Final exam or mini-project: **25%**
- Participation: **5%**

Lecture Format

- The lecture will be offered in an interactive style. We will have two hours of lectures, and one hour of programming tutorials and discussions.

- The lecturer will leverage videos, online courses, and activities to assist the process of teaching.

- The students are encouraged to take an active role in the class. They may participate in the design of assignments and exams.