Syllabus

CS 451
Introduction to Parallel and Distributed Computing
http://www.cs.iit.edu/~iraicu/teaching/CS451-S14/index.html

Semester: Spring 2014
Lecture Time: TBA
Location: TBA

Professor: Dr. Ioan Raicu or Dr. Zhiling Lan
Office Hours Time: TBA
Office Hours Location: TBA

Teaching Assistant: TBA
Office Hours Time: TBA
Office Hours Location: TBA

Course Description
This course covers general introductory concepts in the design and implementation of parallel and distributed systems, covering all the major branches such as Cloud Computing, Grid Computing, Cluster Computing, Supercomputing, and Many-core Computing.

The specific topics that this course will cover are: asynchronous/synchronous computation/communication, concurrency control, fault tolerance, GPU architecture and programming, heterogeneity, interconnection topologies, load balancing, memory consistency model, memory hierarchies, Message passing interface (MPI), MIMD/SIMD, multithreaded programming, parallel algorithms & architectures, parallel I/O, performance analysis and tuning, power, programming models (data parallel, task parallel, process-centric, shared/distributed memory), scalability and performance studies, scheduling, storage systems, synchronization, and tools (Cuda, Swift, Globus, Condor, Amazon AWS, OpenStack, Cilk, gdb, threads, MPICH, OpenMP, Hadoop, FUSE).

Some of these topics are covered in more depth in the graduate courses focusing on specific sub-domains of distributed systems, such as Advanced Operating Systems (CS550), Parallel Computing (CS546), Cloud Computing (CS553), Data-Intensive Computing (CS554), Advanced Computer Architecture (CS570), and Fault Tolerant Computing (CS595).While this CS451 course is not a pre-requisite to any of the graduate level courses in distributed systems, both undergraduate and graduate students who wish to be better prepared for these courses could take this CS451 course. This course involves lectures, programming assignments, and exams. Graduate students who have already taken CS546, CS550, CS553, CS554, CS570, or CS595 should not take this CS495 class. Furthermore, this CS495 class should not be taken concurrently with CS546, CS550, CS553, CS554, CS570, or CS595.

Many of these graduate courses are part of the Master of Computer Science Specialization in Distributed and Cloud Computing. This CS495 course is part of the Undergraduate Specialization in Data Science and the Specialization in Distributed and Cloud computing.

This course was offered as CS495 in the past. Prerequsites: CS351 or CS450.

Textbooks

We will also use be using the textbook Distributed and Cloud Computing: Clusters, Grids, Clouds, and the Future Internet (DCC) by Kai Hwang, Jack Dongarra & Geoffrey C. Fox (Required). This is the most modern book about distributed systems I have found. Some of the fundamental topics in this book are not covered in enough detail, so for some topics, we will use another textbook, Andrew S. Tanenbaum and Maarten van Steen. "Distributed Systems: Principles and Paradigms" (DSPD), Prentice Hall, 2nd Edition, 2007 (Optional). I encourage you to buy both tetxbooks as they are both excellent, but if you have to choose just one, please buy the first (DCC), and the necessary optional reading material needed will be provided to the students in class.

Prerequisites

Systems Programming (CS351) or Operating Systems (CS450)

Mailing lists

There is a course mailing list; you can send mail to the list by sending email to cs495-s14@datasys.cs.iit.edu. Please see http://datasys.cs.iit.edu/mailman/listinfo/cs495-s14 for more information about the course mailing list.

Detailed Course Topics

Lecture topics:

- Distributed System Models
- High-Performance Computing
- Grid Computing
- Cloud Computing
- Many-core Computing
- Many-Task Computing
- Programming Systems and Models
- Processes and threads
- MapReduce
- Workflow Systems
- Virtualization
- Distributed Storage & Filesystems
- Data-Intensive Computing
- Distributed Hash Tables
- Consistency Models
- Fault Tolerance
- Performance analysis and tuning

- Parallel architectures
- Multithreaded programming
- GPU architecture and programming
- Message passing interface

Projects

There will be 6 projects throughout the semester, each worth 10% of the total grade. The projects will be completed in teams of 2 students, and will be programming intensive. Some projects will require knowledge of Java, while others will require knowledge of C and/or C++. It is expected that students know the basics of both of these languages.

Late Policy

Projects will be due at 11:59PM on the day of the due date, through BlackBoard. There will be a 15 minute grace period. Any late submissions beyond the grace period will be penalized 10% every day it is late.

Exams

There will be 2 exams, one covering the material from the first half of the class, and the second covering the material from the second half. The exams will be individual, but students will be allowed to use their textbooks and any notes they have (on paper). No electronic devices such as phones, eReaders, tables, or laptops will be allowed. Simple calculators can be used.

Please note that they extend for 45 minutes after the usual end of class, but this should not interfere with anyone's other classes due to the lunch period.

There will be no makeup exams.

Grades

- **Projects (6): 60%**
- **Exam (2): 40%**

The following grading scale will be used. The scale will be adjusted downwards based on the overall performance of the entire class. Traditionally, in my classes, the class average score will typically fall in the B-grade range.

- **A: 90% ~ 100%**
- **B: 80% ~ 89%**
- **C: 70% ~ 79%**
- **D: 60% ~ 69%**
- **E: 0% ~ 59%**