Syllabus

CS 553: Cloud Computing

http://www.cs.iit.edu/~iraicu/teaching/CS553-F14/index.html

Semester: Fall 2014
Lecture Time: Monday/Wednesday, 11:25AM - 12:40PM
Location: Stuart Building (SB) 104

Professor:

- **Dr. Ioan Raicu** (iraicu@cs.iit.edu, 1-312-567-5704)
  - Office Hours Time: Monday 12:45PM-1:45PM (SB237D), Wednesday 10:20AM-11:20AM (SB237D)

Teaching Assistants (cs553-f14@datasys.cs.iit.edu**):**

- **Ke Wang** (kwang22@hawk.iit.edu)
  - Office Hours Time: Monday/Thursday 3PM-4PM (SB002)
- **Iman Sadooghi** (isadoogh@iit.edu)
  - Office Hours Time: Tuesday/Friday 12:45PM-1:45PM (SB003b)
- **Dongfang Zhao** (dzhao8@hawk.iit.edu)
  - Office Hours Time: Tuesday/Thursday 2PM-3PM (SB002)
- **Tonglin Li** (tli13@iit.edu)
  - Office Hours Time: Wednesday/Friday 2PM-3PM (SB002)

Office Hours Summary:

- **Monday:** 12:45PM-1:45PM (Ioan SB237D), 3PM-4PM (Ke SB002)
- **Tuesday:** 12:45PM-1:45PM (Iman SB003b), 2PM-3PM (Dongfang SB002)
- **Wednesday:** 10:20AM-11:20AM (Ioan SB237D), 2PM-3PM (Tony SB002)
- **Thursday:** 2PM-3PM (Dongfang SB002), 3PM-4PM (Ke SB002)
- **Friday:** 2PM-3PM (Tony SB002), 12:45PM-1:45PM (Iman SB003b)

Course Description

Cloud Computing is "A large-scale distributed computing paradigm that is driven by economies of scale, in which a pool of abstracted, virtualized, dynamically-scalable, managed computing power, storage, platforms, and services are delivered on demand to external customers over the Internet." It has become a driving force for information technology over the past several years, and it is hinting at a future in which we won't compute on local computers, but on centralized facilities operated by third-party compute and storage utilities. Governments, research institutes, and industry leaders are rushing to adopt Cloud Computing to solve their ever-increasing computing and storage problems arising in the Internet Age. There are three main factors contributing to the surge and interests in Cloud Computing: 1) rapid decrease in hardware cost and increase in computing power and storage capacity, and the advent of multi-core architecture and modern supercomputers consisting of hundreds of thousands of cores; 2) the exponentially growing data size in scientific instrumentation/simulation and Internet publishing and archiving; and 3) the wide-spread adoption of Services Computing and Web 2.0 applications. This course is a tour through various topics and technologies related to Cloud Computing. Topics include distributed system models and enabling technologies, computer clusters for scalable Computing, virtual machines and virtualization of clusters and datacenters, design of cloud computing platforms, cloud programming and software environments (Workflow Systems, MapReduce, Google App Engine, Amazon AWS, Microsoft Azure, and emerging cloud software stacks), grid computing and resource management, P2P computing with overlay networks, ubiquitous computing with

clouds and the Internet of things, and data-intensive distributed computing. The course involves lectures, projects, programming assignments, quizzes, and exams. Prerequsites: CS450 or CS455.

Required Texts

We will be using the textbook Distributed and Cloud Computing: Clusters, Grids, Clouds, and the Future Internet by Kai Hwang, Jack Dongarra & Geoffrey C. Fox. (Required)

Prerequisites

CS450 (Operating Systems) or CS455 (Data Communications). Other courses that might contribute to having a better in depth understanding of this course are CS542, CS546, CS451, CS550, CS551, CS570, and CS554; these courses are not required. Many of these graduate courses (including the CS553 course) are part of the Master of Computer Science Specialization in Distributed and Cloud Computing.

Detailed Course Topics

Lecture topics:

- Distributed System Models
- Parallel Computing
- Virtualization
- Cloud Platform Architectures
    - Amazon AWS
    - Microsoft Azure
    - Google App Engine
    - Google MapReduce / Yahoo Hadoop
    - Eucalyptus, Nimbus, OpenStack
- Service-Oriented Architectures
- Cloud Programming
- Grid Computing
- Peer-to-Peer Computing

Programming Assignments

There will be 4 programming assignments throughout the semester, each worth 10% of the total grade, and each taking about 4 weeks to complete. These assignments will be completed in teams of 3 students. The projects will require knowledge of Java, C and/or C++. It is expected that students know the basics of these languages. These assignments must be submitted through BlackBoard, one submission per team.

Computer Usage

Computer systems that can be used for development of projects (more information about access to these will be passed in the first several lectures):

- **Sirius:** 12-node (100-core) private cloud running OpenStack using virtualization and Linux
- **Jarvis:** 10-node (80-cores) Linux Cluster running SGE and NVIDIA GPUs
- **Amazon AWS:** $100 credit per student

Project

There will be 1 mid-term project which will be worth 10% of the grade. This mid-term project will be completed in teams of 3 students. This project will require about 4 weeks to be completed, and will involve in the writeup of a report. This assignment must be submitted through BlackBoard, one submission per team.

Late Policy

Assignments will be due at 11:59PM on the day of the due date, through BlackBoard. There will be a 15 minute grace period. There will also be a 7-day late pass, where students can submit late assignments without penalty; the late pass can be used in 1-day increments spread out over multiple assignments. Any late submissions beyond the grace period and beyond the 7-day late pass, will be penalized 10% every day it is late. Assignments will not be accepted after the solutions have been posted.

Exams & Quizzes

There will be bi-weekly short 15 min quizzes covering material from the prior 2 weeks; each quiz will be worth 5% of the grade. If you are an in-class student, you must take the quiz in class; for remote students, you do not have to go to a testing center to take the quiz, I'll send you an online quiz the day of the quiz which you will have 24 hours to get it done online. The quizzes 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. The quizzes will be on the following schedule:

- **Wednesday, September 10th, 2014 from 12:25PM – 12:40PM in Stuart Building 104**
- **Wednesday, September 24th, 2014 from 12:25PM – 12:40PM in Stuart Building 104**
- **Wednesday, October 8th, 2014 from 12:25PM – 12:40PM in Stuart Building 104**
- **Wednesday, October 22nd, 2014 from 12:25PM – 12:40PM in Stuart Building 104**
- **Wednesday, November 5th, 2014 from 12:25PM – 12:40PM in Stuart Building 104**
- **Wednesday, November 19th, 2014 from 12:25PM – 12:40PM in Stuart Building 104**

There will be 1 final exam that will cover material from the entire semester. 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. The final exam will be worth 25% of the overall grade. In-class students must take the exam in class. Online students can take the exam at an official testing center. The final exam will be scheduled on:

- **Wednesday, December 3rd, 2014 from 11:25AM - 1:25PM in Stuart Building 104**

Please note that the exam is extended for 45 minutes after the usual end of class, but this should not interfere with anyone's other classes due to the lunch period. Also note that the exam is scheduled the week prior to the official final exam week.

There will be no makeup exams or quizzes.

Of the 6 quizzes, I'll drop your lowest quiz score when computing your final grade.

Grades

Grading Policies:

- **Project (1):** 10% -- can use late day passes
- **Programming Assignments (4):** 40% -- can use late day passes
- **Quiz (6):** 25% -- will drop lowest grade
- **Exam (1):** 25%

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 a solid B-grade range.

- **A: 87% ~ 100%**
- **B: 75% ~ 86%**
- **C: 62% ~ 74%**
- **D: 50% ~ 61%**
- **E: 0% ~ 49%**

Mailing lists

This course will use Piazza to facilitate discussions for assignments, at http://piazza.com/iit/fall2014/cs553/home (it has not been setup yet, more instructions will follow). Piazza should be the primary mechanism of communication between the students and the professor and the TAs. If you have a question and want to reach only the TAs and professor, send email to cs553-f14@datasys.cs.iit.edu. As a last resort, send individual emails directly (iraicu@cs.iit.edu, kwang22@hawk.iit.edu, isadoogh@iit.edu, dzhao8@hawk.iit.edu, tli13@iit.edu) when you believe the message is not appropriate to be sent to the entire class, or to all the TAs and professor.