

Syllabus

CS 554: Data-Intensive Computing

<http://www.cs.iit.edu/~iraicu/teaching/CS554-F17/>

Semester: Fall 2017

Lecture Time: Monday/Wednesday, 1:50PM - 3:05PM

Location: Robert A. Pritzker Science Ctr 111

Professor:

- **Dr. Ioan Raicu** (iraicu@cs.iit.edu, 1-312-567-5704)
 - Office Hours: Wednesdays 3:10PM-4:10PM (SB237D)

Teaching Assistant:

- **Poornima Nookala** (pnookala@hawk.iit.edu)
 - Office Hours: Monday/Wednesdays 11:30AM-12:30PM (SB007)



Ioan Raicu



Poornima Nookala

Overview

This course is a tour through various research topics in distributed data-intensive computing, covering topics in cluster computing, grid computing, supercomputing, and cloud computing. We will explore solutions and learn design principles for building large network-based computational systems to support data intensive computing. This course is geared for junior/senior level undergraduates and graduate students in computer science.

Course Goals

The support for Data Intensive Computing is critical to advancing modern science as storage systems have experienced an increasing gap between its capacity and its bandwidth by more than 10-fold over the last decade. There is an emerging need for advanced techniques to manipulate, visualize and interpret large datasets. Building large scale distributed systems that support data-intensive computing involves challenges at multiple levels, from the network (e.g., transport, routing) to the algorithmic (e.g., data distribution, resource management) and even the social (e.g., incentives). This course is a tour through various research topics in distributed systems, covering topics in cluster computing, grid computing, supercomputing, and cloud computing. We will explore solutions and learn design principles for building large network-based computational systems to support data intensive computing. Our readings and discussions will help us identify research problems and understand methods and general approaches to design, implement, and evaluate distributed systems to support data intensive computing. Topics include resource management (e.g. discovery, allocation, compute models, data models, data locality, virtualization, monitoring, provenance), programming models, application models, and system characterization. Our discussions will often be grounded in the context of deployed distributed systems, such as the XSEDE, Amazon EC2 and S3, various top supercomputers (e.g. IBM BlueGene/Q, Cray XK7), various cloud platforms (e.g. Amazon AWS, Google App Engine, Windows Azure, and private clouds), various software/programming platforms (e.g. Hadoop/MapReduce, Dryad, Spark/Sparrow/Mesos, Swift/Falcon, MATRIX, and Parrot/Chirp), and various file systems (e.g. GPFS, Lustre, PVFS, FusionFS, and Ceph). The course involves lectures, outside invited speakers, discussions of research papers, and a major project (including written reports and oral presentations).

Required Texts

None; readings will be from published research online material.

Prerequisites

[CS450](#) (Operating Systems). Other courses that might contribute to having a better in depth understanding of this course are [CS542](#), [CS546](#), [CS451](#), [CS550](#), [CS551](#), CS552, [CS553](#), [CS570](#), and/or CS595 (VMs). These courses are not required, and the introductory content of some of those courses will be repeated in this class. Many of these graduate courses are part of the [Master of Computer Science Specialization in Distributed and Cloud Computing](#).

Knowledge of the following topics is recommended to get the most out of this class:

1. Programming (C, C++, or Java)
2. Networking
3. Operating systems
4. Architecture
5. Distributed systems

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):

- **Amazon AWS** (<https://aws.amazon.com>): To be used for Programming Assignment #3 (likely \$100 per student)
- **Chameleon** (<https://www.chameleoncloud.org>): To be used for Programming Assignment #1 and #2 (likely 1000 CPU hours per student)

Detailed Course Topics

Lecture topics:

- Paradigms
 - Supercomputing (e.g. IBM BlueGene/P/Q, Cray XK7)
 - Grid Computing (e.g. XSEDE, OSG)
 - Cloud Computing (e.g. Amazon AWS, Google App Engine, Windows Azure)
 - Many-core Computing (e.g. NVIDIA GPUs, Xeon Phi)
- Parallel Programming Systems
 - MapReduce (e.g. Hadoop, Spark)
 - Workflows (e.g. Swift)
 - MPI (e.g. MPICH)
 - OpenMP/Cilk/CUDA/PThreads
- Job Management Systems
 - Batch scheduling (e.g. Condor, Slurm, SGE, PBS)
 - Light-weight Task Scheduling (e.g. Falkon, Sparrow, MATRIX, CloudKon, Albatross, Spark)
- Storage Systems
 - Shared File Systems (e.g. NFS)
 - Distributed File Systems (e.g. HDFS, FusionFS, Ceph)
 - Parallel File Systems (e.g. GPFS, PVFS, Lustre)
 - Distributed NoSQL Key/Value Stores (e.g. Casandra, MongoDB, ZHT, DynamoDB)
 - Relational Databases (e.g. MySQL, HRDBMS)

Projects

There will be a major semester long project that will require the implementation of a real/simulated system, a written report, and an oral presentation. Projects can fall in a number of different areas, which are related to data-intensive distributed computing. These areas will be expanded with specific projects in each area.

- Storage systems
- Scheduling
- Run-time systems
- Programming systems

The students will be given several (3~10) project descriptions (<1 page in length for each project), and students must pick 1 project to develop further into a proposal (2~3 pages). Students are encouraged to work in teams of 2 or 3 on these projects, and multiple teams can work on the same project (independently). These projects will generally be of high interest to the [DataSys Laboratory](#) and are likely sub-projects from existing [ongoing research projects](#). These projects could have some interaction with PhD students in the DataSys lab who are leading the projects at large. Students who obtain excellent results on their projects will be encouraged to submit their work as a publication to highly respected conferences and workshops.

Some major projects and tools you should familiarize yourself with are (as they might be useful to your projects):

- Operating systems: Linux
- Scripting: BASH
- Source control: SVN, GitHub
- Programming languages: Java, C/C++, Python
- Job submission systems: GRAM, PBS, Condor, Cobalt, SGE
- Programming models: MapReduce (Hadoop/Spark), MPI (MPICH), Multi-Threading (PThreads/CUDA/OpenMP/Cilk), Workflows (Swift)
- File systems: FUSE
- Parallel file systems: GPFS, PVFS, Lustre
- Distributed file systems: GPS, HDFS, FusionFS, Ceph, GlusterFS
- Data services: GridFTP
- Grid middleware: Globus
- Cloud middleware: Nimbus, Eucalyptus, OpenNebula, OpenStack
- Key/Value Stores: Chord, Tapestry, ZHT, Cassandra, MongoDB, MemCached, DynamoDB
- Message Queues: SQS, Kafka, Fabriq
- Simulation environments: GridSim, SimGrid, OptorSim, GangSim, Bricks, SimMatrix, PeerSim, CODES/ROSS
- Virtualization: Oracle Virtual Box, XEN, VMWare

Schedule

Tentative schedule:

- 08/21/2017: Syllabus
- 08/23/2017: Introduction to Distributed Systems
- 08/28/2017: Introduction to Distributed Systems
- 09/01/2017: Last day to add/drop class
- 09/04/2017: NO CLASS (Labor Day)
- 09/06/2017: Project Brainstorming
- 09/13/2017: **Group formation due; Project Proposal Due**
- 10/18/2017: **Project Midterm Progress Report Due**
- 11/27/2017: **Project Final Presentations**
- 11/29/2017: **Project Final Presentations**
- 12/01/2017: **Project Final Reports Due**
- 12/13/2017: Grades Due

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 4-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. Only project related assignments can use the late day passes. Paper reviews you write for reading assignments cannot use late day passes, and will incur a 20% penalty per day for late submissions. Any eligible late submissions (e.g. project proposal, project mid-term report, and project final report) beyond the grace period and beyond the 4-day late pass, will be penalized 20% every day it is late.

Mailing lists

This course will use Piazza to facilitate discussions around the projects, at <http://piazza.com/iit/fall2017/cs554/home>.

Participation

There will be 20% points of the overall grade that will be received for participation. Students will be assigned research papers throughout the semester that they need to lead a discussion for. Online students can participate through pre-recorded videos. Paper reviews will also go towards participation points.

Grades

There are no exams or homework.

Grading Policies:

- Project Proposal: 15%
- Mid-semester Progress Report: 15%
- Final Oral Presentation: 25%
- Final Project Report: 25%
- Participation: 20%

The following grading scale will be used. The scale will be adjusted downwards based on the overall performance of the entire class, if the class average score falls below the B-grade range.

- **A: 87% ~ 100%**
- **B: 75% ~ 86%**
- **C: 60% ~ 74%**
- **E: 0% ~ 59%**

Course Outcomes

When a student completes this course, s/he should be able to:

1. Understand the importance of data-intensive computing
2. Understand the difference between cluster, grid, clouds, and supercomputing.
3. Understand how to build large scale distributed systems
4. Understand applications that require data-intensive computing
5. Understand trends in many-core computing and challenges that will come with them
6. Build distributed systems
7. Be familiar with multiple programming models
8. Read and understand systems research papers
9. Make a formal presentation on a technical topic
10. Write up a formal report (or even a research paper) on the project