# Syllabus

### *EECS 395 / EECS 495*

### *Hot Topics in Distributed Systems: Data-Intensive Computing*

http://www.eecs.northwestern.edu/~iraicu/teaching/EECS495-DIC/index.html

**Quarter: Winter 2010**

**Time: Tuesday/Thursday, 12:30PM - 1:50PM**

**Location: TECH L158**

**Instructor: Dr. Ioan Raicu (iraicu@eecs.northwestern.edu)**

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 TeraGrid, Amazon EC2 and S3, various top supercomputers (e.g. IBM BlueGene/P, Sun Constellation, Cray XT5), and various software/programming platforms (e.g. Google's MapReduce, Hadoop, Dryad, Sphere/Sector, Swift/Falkon, and Parrot/Chirp). The course involves lectures, outside invited speakers, discussions of research papers, and a major project (including both a written report and an oral presentation).

## Required Texts

None; readings will be from published research online material.

## Prerequisites

Graduate students, none; undergraduates, EECS110, EECS111, EECS211, EECS311, EECS 340, EECS 343, EECS345; students please contact the instructor to ensure proper background.

Prerequisites by topics:

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

## Detailed Course Topics

Lecture topics:

- Distributed Systems: Clusters, Supercomputers, Grids and Clouds
- Data Intensive Computing Overview
- Local Resource Management Systems
- Storage Systems
- Shared, Distributed and Parallel File Systems
- Parallel I/O
- Scientific Computing and Applications
- Parallel Programming Systems and Models
- MapReduce & Hadoop
- Sphere/Sector
- Parrot and Chirp
- Swift/Falkon
- Data-Intensive Computing with GPUs
- Data-Intensive Computing with Databases
- Many-core Computing Era and New Challenges
- Open Research Questions in Data-Intensive Computing

## Computer Usage

None required. Computer systems that can be used for development of projects and assignments:

- **falkon.eecs.northwestern.edu**
    - request account by sending email to iraicu@eecs.northwestern.edu
    - Intel Xeon, 16-cores @ 2.33GHz, 48GB RAM, 7TB RAID5 disk, 1Gb/s network
        - Primary: Linux Suse 11.2 x64
        - Virtual Machine: Windows Server 2008 x64
    - AMD Atholon II X4, 4-cores @ 2.6GHz, Nvidia GTX295 with 2GB RAM and 800 cores, 4GB RAM, 75GB disk, 1Gb/s network
        - Primary: Windows 7 x64
        - Virtual Machine: Linux SuSe 11.2 x64

Other systems that can be used:

- PADS Cluster at University of Chicago (1K cores x64)
- IBM BlueGene/P at Argonne National Laboratory (160K PPC)
- SiCortex at Argonne National Laboratory (5832 MIPS)
- ANL/UC TG Cluster at Argonne National Laboratory (~200 IA32)
- TeraGrid (150K of all variety of CPUs)

- Sun Constellation at TACC (62K x64)
- Magellan at Argonne National Laboratory (10K x64 cloud)
- Amazon EC2

## Laboratory Projects

There will be a major quarter long project (on the topic of choice of the student) that will require possibly 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 somewhat related to data-intensive distributed computing. These areas will be expanded on during the course, with specific projects for each area.

- Distributed file systems
- Data aware scheduling algorithms
- Distributed operating systems
- Distributed job management systems
- Parallel programming languages
- Distributed workflow systems
- Distributed monitoring systems
- Scientific computing with GPUs
- Scientific computing with MapReduce
- Distributed caching strategies
- Distributed cache eviction policies
- Distributed hash tables
- Virtualization impact for data-intensive computing

More project ideas from the Globus community can be found at http://dev.globus.org/wiki/Project_Ideas; not all these projects are suitable for this course, but they will help your brainstorming exercise in finding a suitable topic.

Main operating systems, software components, and major projects you should familiarize yourself with are (they might be useful to your projects):

- Operating systems: Linux, Windows
- Scripting: BASH
- Source control: SVN
- Programming languages: Java, C/C++
- Job submission systems: GRAM, PBS, Condor, Cobalt, SGE, Falkon
- Programming models: MapReduce (Hadoop), MPI (MPICH), Multi-Threading (PThreads), Workflows (Swift, Pegasus/DAGMan, Nimrod, Taverna, BPEL)
- File systems: FUSE
- Parallel file systems: GPFS, PVFS, Lustre
- Distributed file systems: GPS, HDFS
- Data services: GridFTP
- Grid middleware: Globus
- Cloud middleware: Nimbus, Eucalyptus, OpenNebula
- Distributed hash tables: Chord, Tapestry
- Simulation environments: GridSim, SimGrid, OptorSim, GangSim, Bricks
- Virtualization: Sun Virtual Box, XEN, VMWare

## Mailing lists

The course mailing list can be found at [http://www.eecs.northwestern.edu/mailman/listinfo/eecs495-dic](http://www.eecs.northwestern.edu/mailman/listinfo/eecs495-dic). If you are registered for the course, you have been already signed up for it, but otherwise, please sign up.

## Grades

Grading Policies:

- Participation in paper discussions (including writeups for papers): 15%
- Homeworks (~5 assignments): 20%
- Mid-quarter oral presentation: 5%
- Final oral presentation: 10%
- Project / Report: 50%

## 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 computing, grid computing, supercomputing, and cloud computing
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 a research paper
9. Make a formal presentation on a technical topic
10. Write up a formal report (or even a research paper) on the project