New York University Tandon School of Engineering
Department of Computer Science & Engineering

CS 9223 Fall 2018
Distributed Systems Professor Sandoval

Large-scale distributed systems are the core software infrastructure underlying cloud computing. These systems consist of tens of thousands of networked computers working together to provide unprecedented performance and fault-tolerance. Examples include distributed databases (e.g. Google's Spanner, Amazon's S3 and Dynamo), distributed computation frameworks (e.g. Hadoop/Spark and TensorFlow) or even decentralized currency systems (e.g. BitCoin). This class teaches the abstractions, design and implementation techniques that allow you to understand and build such distributed systems.

It will present abstractions and implementation techniques for engineering distributed systems. Major topics include multithreading, network programming, **fault tolerance, replication, and consistency**. We will also spend a lot of time studying several case studies of distributed systems.

Contact

Email: gustavo.sandoval@nyu.edu
2 Metrotech, 10th Floor, Office 10.040
Office hours: Wed: 2:30 – 4:00

Student Assistants & Office Hours:

- Adam Sanghera – as8513@nyu.edu
- Dov Salomon – dsalomon@nyu.edu

Prerequisites:

CS 6233 – Intro to Operating Systems.
Expertise in C/C++ programming.

Academic Honesty

All work submitted in this course must be your own. Cheating and plagiarism **will not be tolerated**. If you have any questions about a specific case, *please ask me.*
NYU Poly's Policy on Academic Misconduct is here:

http://engineering.nyu.edu/academics/code-of-conduct/academic-misconduct

Course requirements

- This term we will be using Piazza for class discussion. The system is highly catered to getting you help fast and efficiently from classmates, the TA and myself. Rather than emailing questions to the teaching staff, I encourage you to post your questions on Piazza. Here's the link for our class: https://piazza.com/nyu/fall2018/cs9223/home
- In class quizzes and participation in Piazza will count towards 10% of your grade. So Participate!!
-

Course schedule (***Tentative and subject to change***)

Week	Date	Topic	Reading
1	**Sept 6**	Introduction to Distributed Systems and : MapReduce	MapReduce
2	**Sept 13**	Go Intro. Web Architectures and Cloud Computing	Online go Tutorial Guest Speaker: Andy Cheung (SignalSciences)
3	**Sept 20**	Distributed Communication: Sockets and RPC	Online go Tutorial
4	**Sept 27**	Time, Clocks and Distributed Systems	Logical Clocks
5	**Oct 4**	Fault Tolerance: Primary/backup Replication	GFS VR Revisited
6	**Oct 11**	CAP Theorem Consistency Models	Perspectives on the CAP Theorem
7	**Oct 18**	Consensus: Paxos	Paxos Made Simple
8	**Oct 25**	Consensus: Raft	Raft
9	**Nov 1**	Consensus in Practice**:** Chubby, and BigTable	Special Guests from Google to talk about: Chubby, BigTable
10	**Nov 8**	Relaxed Consistency	Dynamo Bayou
11	**Nov 15**	Distributed Transactions	Spanner
	Nov 22	Thanksgiving	

12	**Nov 29**	Peer to peer: Bitcoin	Bitcoin Summary
13	**Dec 6**	Project Presentations	
14	**Dec 13**	Project Presentations	

Textbook: None.

Useful Reference Books:

- Designing Data Intensive Applications (Kleppmann)
- Distributed Systems (Tanenbaum and Steen)
- Distributed Systems (Colouris, et al)
- Distributed Systems for Fun and Profit

Reading

We will be reading published papers and discussing them in class. The papers will also be covered in the exams.

Acknowledgements

This course builds upon several existing distributed systems courses from other universities:

MIT Distributed Systems course by Robert Morris.

NYU Distributed Systems course by Jinyang Li

Columbia Distributed Systems course by Roxana Geambasu

Grading

Grading will be based on the following weights.

30 % Assignments
30 % Quizzes
40 % Project

Assignments:

There will be several assignments involving programming projects. Every assignment has to be done individually by each student. Discussions between students and help in using the programming environment are permitted and encouraged, however no code or solutions may be copied. The assignments must be in Completed in **Go**. There is no late policy for assignments; any absence notifications should be communicated via http://engineering.nyu.edu/life/student-affairs/advocacy-privacy-and-compliance.

Project:

The goal of this project is for the student to gain experience in implementing a Distributed System different to what we see in class. The project will be performed in teams, details to be announced. Students will be responsible for selecting a domain area problem (some example problems will be suggested). The grade for the project will be divided between: problem and background research, proposed solutions (short presentation), and implementation. More details will be provided in class.

Quizzes:

There will be a series of very short quizzes at the beginning of selected classes covering material to date. The lowest quiz mark will be dropped. Quiz marks will not be excused.

Other Grading notes:

Please take the following into consideration during and after the semester and save yourself one or many emails.

1) **I must grade every student EXACTLY the same way.** To this end, I cannot give you special consideration as a result of your academic status (probation or otherwise), scholarships, work status, family situation, visa status, race, color, creed, religious beliefs, past alien abductions, current moon cycle, location of the sun in the sky or anything other than your academic performance. **Your grade must be based on your academic performance in my class**.
2) **I cannot change your grade simply because you ask me to.** Your grade is calculated based on your performance from the first day of class to moment you turn in the final exam.
3) **I will not give you additional work.** Please remember that I must treat all students the same, so if I give you additional work, I would have to give it to the entire class. This is unfair to the students who complete their work on time.
4) **Your grade is a measure of your performance in my class.** If you receive an "F" it is because you have demonstrated that you do not understand the material in the course; if you receive an "A" it is because you have demonstrated that you fully understand the material covered in the course. Other grades are assigned accordingly.

Moses Center Statement of Disability

If you are student with a disability who is requesting accommodations, please contact New York University's Moses Center for Students with Disabilities at 212-998-4980 or mosescsd@nyu.edu. You must be registered with CSD to receive accommodations. Information about the Moses Center can be found at www.nyu.edu/csd. The Moses Center is located at 726 Broadway on the 2nd floor.