18-847F: Special Topics in Computer Systems

Foundations of Cloud and Machine Learning Infrastructure

Lecture 1: Logistics and Overview

Foundations of Cloud and Machine Learning Infrastructure

Graduate Seminar Class

Few Lectures

Reading research papers

Student presentations

Class Discussions

Final Research Project (No Exams!)

Learning Objectives

- Know the state-of-the-art frameworks in cloud and machine learning and their theoretical foundations
- Read and provide constructive criticism of research papers
- Present to an audience, and answer their questions
- Do creative, collaborate research

Why study Cloud and ML infrastructure?

What are the largest words after 'Big Data'?

Big Data Gold Rush

Who got rich in the California gold rush?

Big Data Gold Rush

Who got rich in the California gold rush?

Google Compute Engine

amazon web services

IBM Bluemix™

In the Big Data rush, it's the infrastructure companies

Topics Covered

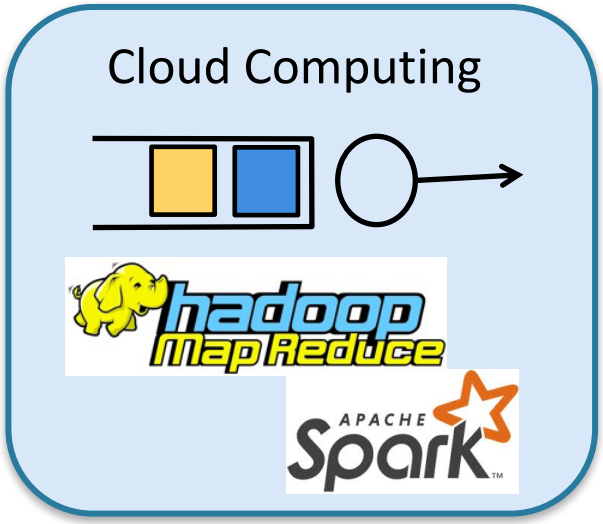

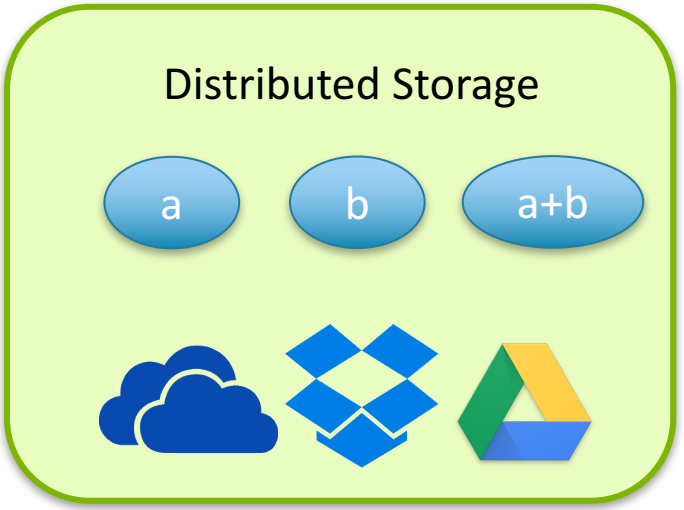

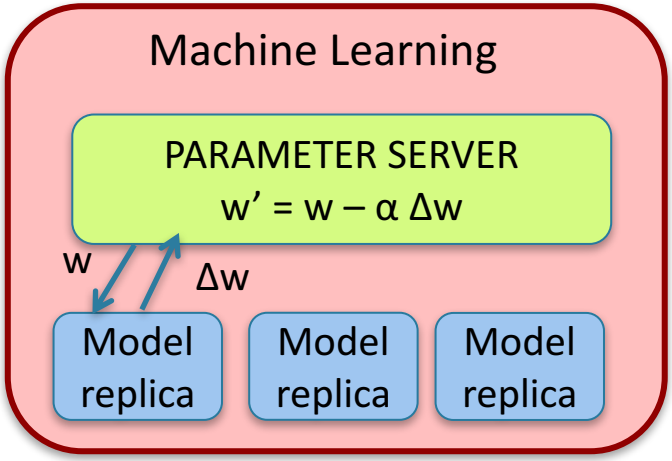

Topics Covered

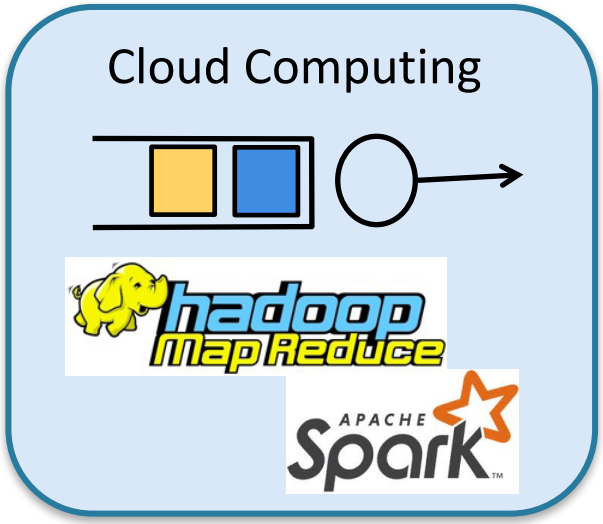

- Scheduling in Parallel Computing
  - MapReduce, Spark
  - Straggler Replication
- Task Replication in Queueing Systems

Topics Covered

- Coding for locality/repair
- Systems implementation of codes
- Reducing latency in content download

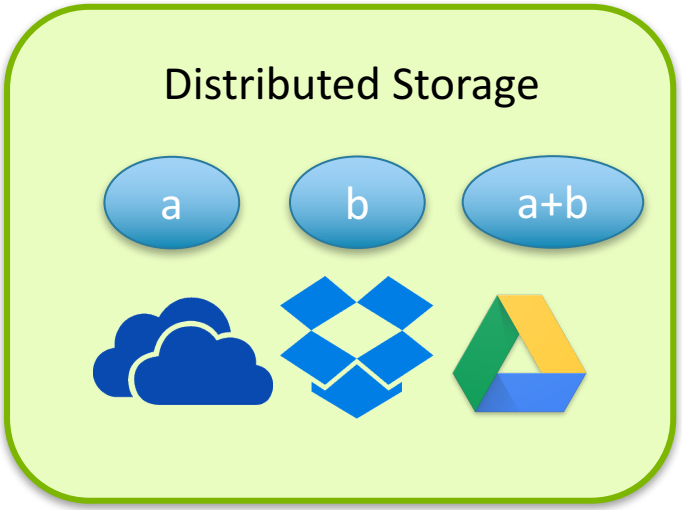

Topics Covered

- SGD and its convergence
- Distributed Deep Learning
- Hyper-parameter tuning
- GANs, Deep reinforcement learning

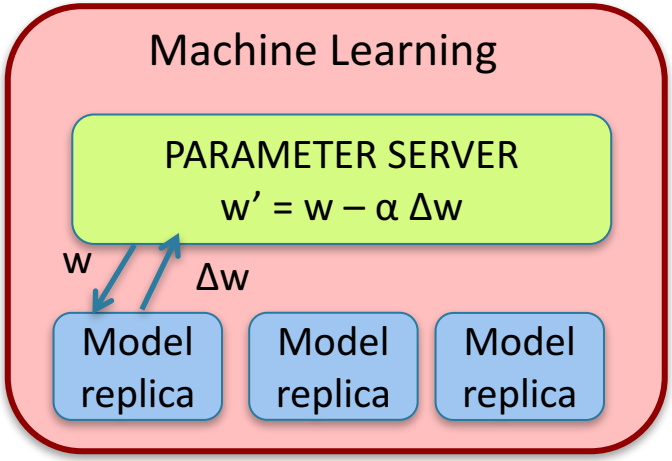

Instructor: Gauri Joshi

B.Tech+M.Tech
2005-2010

SM + PhD
2010-2016

Research Staff Member
2016-2017

Carnegie Mellon University

Assistant Professor
Fall 2017 -

Internships

Bell Labs 

Have worked in all these areas

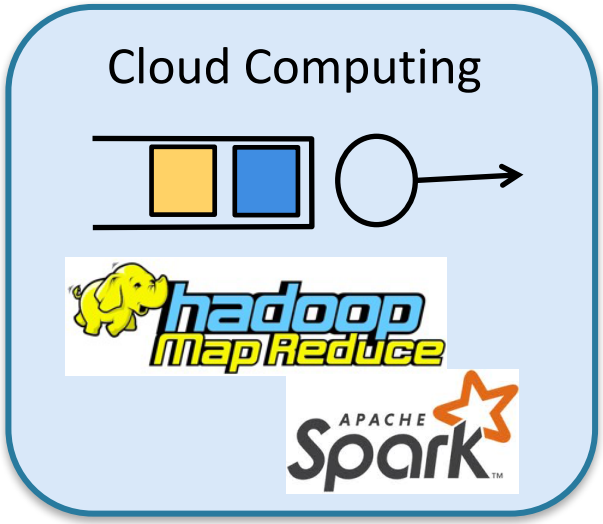

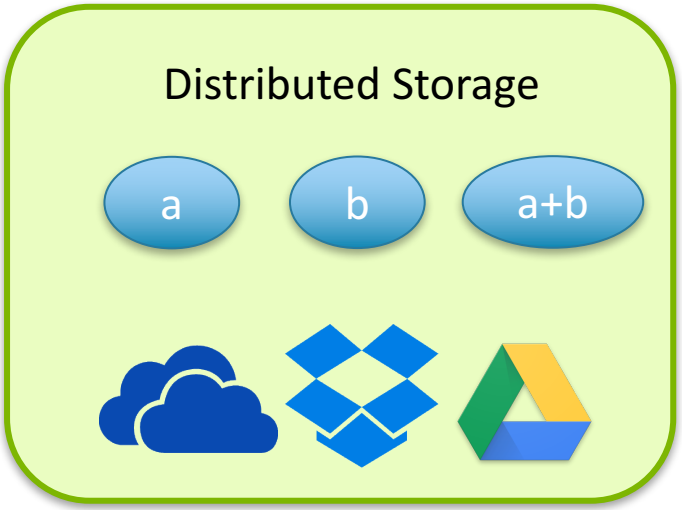

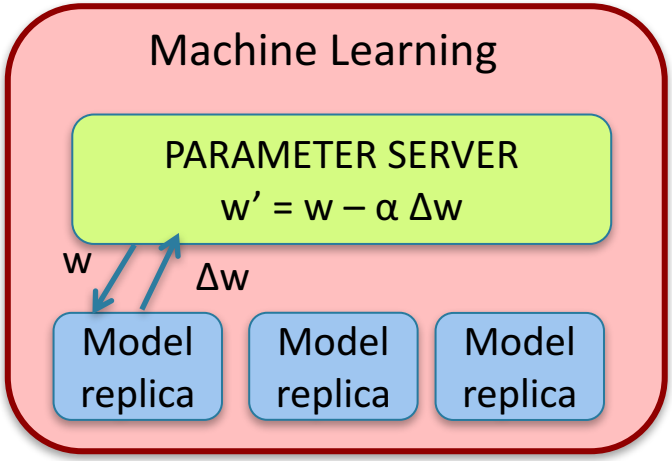

Student Introductions

- Name?
- Department?
- Undergrad/Masters/PhD?
- Previous related classes (if any)?
- What you are looking to learn from this class?

Waiting list will be cleared soon!

Class Hours and Website(s)

- When: Mon, Wed 4:30-6:00 pm
- Where: Scaife Hall 222

- Class Website (Readings, Schedule): https://www.andrew.cmu.edu/user/gaurij/18-847F-Fall-2018.html
- Canvas Site (Readings, Assignments, Projects): https://canvas.cmu.edu/

- No prerequisites. Basic knowledge of probability and linear algebra is encouraged.

Reading Material

Papers will be posted on the class website or on Canvas

- Book chapters
- Survey papers
- Theory papers (Scheduling, Queuing, Coding, Optimization)
- Systems papers (Cloud, Machine Learning)

Additional reference books listed in the syllabus

Instructor/TA and Office Hours

Instructor: Prof. Gauri Joshi (gaurij [AT]andrew.cmu.edu)

TA: Jianyu Wang (jianyuw1 [AT]andrew.cmu.edu)

Office Location: CIC 4105

Office Hours: By appointment

Graduate Seminar Class

A few lectures

Reading research papers

Student presentations

Class Discussions

Final Research Project

Lectures

- Next week: Deeper Overview of probability and queuing theory

- Guest lectures during the semester by authors of papers relevant to this class

Graduate Seminar Class

A few lectures

Reading research papers

Student presentations

Class Discussions

Final Research Project

Homeworks (~50%)

- Submit paper review (due 10:00 am before class)
  - ~Two reviews per week
- Discussion with classmates is okay, but write reviews in your own words.

Paper Review Format

- Summary of the paper
  - Reflects your understanding of the paper
  - Significance & correctness of results
- Discussion Questions for Class (at least 2)
  - Confusions about the paper, open research directions
- Answers to concept-check questions

Homework Grading Rubric (Total: 10 pts)

- Understanding of the paper (4 pts)
- Discussion Questions (3 pts)
- Concept-check questions (3 pts)

Graduate Seminar Class

A few lectures

Reading research papers

Student presentations

Class Discussions

Final Research Project

Class Presentations (~15%)

- Sign up for presentation at least 1 week in advance
- Each student will present 1-2 times in the semester (depends on # of students registered)
- 20 min presentation, followed by 25 min discussion
  - Motivation and Related work
  - Summary of main results
  - Your views on the paper

Presentation Grading Rubric (Total: 10 pts)

- Motivation (1.5 pts)
- Clarity (1.5 pts)
- Understanding/Correctness (4 pts)
- Peer-review Feedback (3 pts)

Graduate Seminar Class

A few lectures

Reading research papers

Student presentations

Class Discussions

Final Research Project

Class Participation (~15%)

- The class will be divided into groups of 3-4 students each
- Each group will discuss one of the discussion questions among themselves
- Summarize the discussion to the whole class

Participation Grading Rubric (Total: 5 pts)

- Attendance and attention (1.5 pt)
- Speaking up in class (1.5 pt)
- Insightful Questions/Comments (2 pt)

Graduate Seminar Class

A few lectures

Reading research papers

Student presentations

Class Discussions

Final Research Project

Research Project (~20%)

- Groups of 1-3
- Original research on a topic of your choice
  - Topics aligned with your research allowed and encouraged
  - If you can't think of topics, come talk to Jianyu or me
- Possible Project Types:
  - New theoretical analysis
  - Implementation using one of the frameworks discussed
  - In-depth literature survey of a particular topic

Timeline

- 1-page proposal due Oct 3
- Publishable quality report (max 5 pg) in ACM format
  - Initial draft due: Nov 21
  - Final report due: Dec 7
- Last week of class: Presentations (~25 min per group)
- Peer-review other presentations

Project Grading Rubric (Total: 20 pts)

- Originality (1 pts)
- Review of Related Work (1.5 pts)
- Writing and Organization (1.5 pts)
- Technical Results (4 pts)
- Final presentation (10 pts)
- Peer-Review (2 pts)

In Summary..

- Paper Reading
- Submitting Reviews
- Class Presentations
- Final Project

Might seem like a lot of work but..

- You will get fast and efficient at reading papers
- The project will be a fun, collaborative exercise
- No exams!

TO DO

- Fill out the sign-up sheet
- Sign-up for presentations
- Start reading the papers
- Form groups for class projects
- Start thinking about projects