

CS 3410: Distributed Systems

Spring 2018	Paper (due Thursday)	Topic	Project
Jan 8-12	1. Google File System	the internet	Go tour
Jan 15-19 (<i>MLK Day</i>)	2. Bigtable	Go, RPC	
Jan 22-26	3. Case study: Google	peer-to-peer	Go exercises
Jan 29-Feb 2	4. Chord	concurrency	Chat service
Feb 5-9	5. Dynamo		Chord: linked list ring
Feb 12-16	6. Case study: Facebook		Chord: finger tables
Feb 19-23 (<i>Presidents' Day</i>)	7. Borg	containers	Chord: fault tolerance
Feb 26-Mar 2	8. MapReduce		1st presentation
Mar 5-9	9. Paxos	consensus	
Mar 12-16 (<i>Spring break</i>)	—	—	—
Mar 19-23	10. Chubby		
Mar 26-30	11. Case study: Twitter		
Apr 2-6	12. Dremel/BigQuery	databases	Paxos
Apr 9-13	13. Megastore		
Apr 16-20	14. Spanner		
Apr 23-27 (<i>Wednesday last day</i>)	—		MapReduce

Resources

- [Syllabus](#)
- [Examples from class](#)
- [A Tour of Go \(start here to learn Go\)](#)
- [Effective Go](#)
- Recommended book: [The Go Programming Language](#)
- [Go package docs](#)
- [RPC Chapter from *Network Programming with Go*](#)
- [Mandelbrot zoom mapreduce package](#)
- Vocabulary:
 - ACID transactions
 - Serializable transactions
 - CAP theorem
 - At least once, at most once, exactly once
 - Idempotent
 - 2-phase commit
 - Pessimistic transactions, optimistic transactions

Projects

- Project 0: [A tour of Go](#)
- Project 1: Go exercises (in CodeGrinder)
- Project 2: [RPC chat room](#) [[plain](#)] [[pdf](#)]
- Project 3: [Chord key/value service](#) [[plain](#)] [[pdf](#)]
- Project 4: [Paxos key/value service](#) [[plain](#)] [[pdf](#)]
- Project 5: [MapReduce](#) [[plain](#)] [[pdf](#)]
- First presentation
- Second presentation

Code to discover your own IP address. This does not work in all cases, but it is a useful starting point:

- [getLocalAddress](#)

Papers

1. [The Google File System](#)
2. [Bigtable: A Distributed Storage System for Structured Data](#)
3. Case study: Google
 - [Web search for a planet: The Google Cluster Architecture](#)
 - [Building Software Systems at Google and Lessons Learned \(video 1:22:44\)](#)
4. [Chord: A Scalable Peer-to-peer Lookup Service for Internet Applications](#)
5. [Dynamo: Amazon's Highly-available Key-value Store](#)
6. Case study: Facebook
 - [Scaling Facebook to 500 Million Users and Beyond](#) (high-level overview, lessons learned)
 - [Scale at Facebook](#) (video, 1 hour)
 - [Needle in a haystack: efficient storage of billions of photos](#) (details about one specific service)
7. [Large-scale cluster management at Google with Borg](#)
8. [MapReduce: Simplified Data Processing on Large Clusters](#)
9. Paxos
 - [The Part-Time Parliament](#)
 - [Paxos Made Simple](#)
 - [Paxos in 25 lines](#)
10. [The Chubby lock service for loosely-coupled distributed systems](#)
11. Case study: Twitter
 - [Real-Time Delivery Architecture at Twitter](#) (56 minute video)
 - [Decomposing Twitter: Adventures in Service-Oriented Architecture](#) (50 minute video)
12. Dremel and BigQuery
 - [Dremel: Interactive Analysis of Web-Scale Datasets](#)
 - [An Inside Look at Google BigQuery](#)
 - [BigQuery under the hood](#)
13. [Megastore: Providing Scalable, Highly Available Storage for Interactive Services](#)
14. [Spanner: Google's Globally-Distributed Database](#)

Paper summaries

For each paper that we read, you must write a brief summary and submit it using Canvas. You should write about 1000 words. Here are a few guidelines:

- Be terse and concise. Avoid adjectives and other flowery language. You have very little space to summarize a lot of information.
- Summarize the technical content of the reading. Give equal weight to each section of what you read. If you read 10 pages, you should devote roughly 100 words to each page.
- If you do not understand something, try searching for additional information on Wikipedia or other sources on the web (there is a useful tool called Google...) or come to my office hours. Note that this works best if you do the reading early.

Things to avoid:

- Do not write about your reaction to the reading. Do not waste space telling if you liked it or did not like it or if you found it confusing or you were surprised by it.
- Do not write much about the background material and context. Focus on the specific question posed for each reading.
- Do not ask questions in your summary. The time for questions is before you write it.
- Do not use the first person. You are writing about the technical content of the reading assignment, not about yourself.
- Avoid the passive voice.

If you are unsure about any of these guidelines, talk to me, or take your answer to the writing center to get help from a writing tutor.

Presentations

- [The Byzantine Generals Problem](#)
- [Scale and Performance in a Distributed File System \(AFS\)](#)
- [Large-scale Incremental Processing Using Distributed Transactions and Notifications](#)
- [Time, Clocks, and the Ordering of Events in a Distributed System](#)

- [Paxos Made Live—An Engineering Perspective](#)
- [Life beyond Distributed Transactions: an Apostate's Opinion](#)
- [There is More Consensus in Egalitarian Parliments](#)
- [PNUTS: Yahoo!'s hosted data serving platform](#)
- [A Note on Distributed Computing](#)
- [In Search of an Understandable Consensus Algorithm](#)
- [Mesa: Geo-Replicated, Near Real-Time, Scalable Data Warehousing](#)
- [High-Availability at Massive Scale: Building Google's Data Infrastructure for Ads](#)
- [Distributed Snapshots: Determining Global States of Distributed Systems](#)
- [Principles of Robust Timing over the Internet](#)
- [Twitter Heron: Stream Processing at Scale](#)
- [CAP Twelve Years Later: How the "Rules" Have Changed](#)
- [F1: A Distributed SQL Database That Scales](#)
- [Flexible Paxos: Quorum intersection revisited](#)