CS 3410: Distributed Systems

Spring 2022	Topic	Project (due Wednesday)	Paper (due Friday)
Jan 10–14	Go, RPC	Go basics 1–3	1. Google File System
Jan 17–21 (*MLK Day*)	Go examples	Go basics 4–7	2. Bigtable
Jan 24–28	Effective Go, replicated state machines	Go basics 8–12	3. Paxos
Jan 31–Feb 4	TCP, sockets, clusters	Paxos: prepare	4. Case study: Google
Feb 7–11	coherent caching, CAP	Paxos: accept and decide	5. Chubby
Feb 14–18	transactions, 2-phase commit	MUD command parser	6. Megastore
Feb 21–25 (*Presidents' Day*)	time, clocks, snapshots	MUD single player	7. Spanner
Feb 28–Mar 4	peer to peer	MUD multi player	8. Chord
Mar 7–11	concurrency, actors	Chord: linked list ring	9. Case study: Facebook
Mar 14–18 (*Spring break*)	—	—	—
Mar 21–25	databases	Chord: full chord protocol	10. Calvin
Mar 28–Apr 1	big data		11. MapReduce
Apr 4–8	SOA, microservices	MapReduce: database handling	12. Case study: Twitter
Apr 11–15	eventual consistency	MapReduce: map and reduce	13. Dynamo
Apr 18–22		MapReduce: master node	14. RDDs (Spark)
Apr 25–29 (*Wednesday last day*)			—

Changes to the schedule will be announced in class.

Resources

- Syllabus
- Examples from class
- Effective Go
- Recommended book: The Go Programming Language
- Go package docs
- Screencast on setting up Go and vim-go
- TCP videos
  - TCP service model (16:27)
  - The end-to-end principle (10:33)
  - Sliding window (19:25)
  - Retransmission strategies (9:45)
- RPC demo app in Go
  1. introduction (8:22)
  2. server RPC (3:01)
  3. client RPC (4:50)
  4. command-line flags (13:58)
  5. call function (6:58)
  6. client shell (14:35)
  7. actor (15:23)

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. Paxos
    - Paxos Made Simple
    - Paxos vs Raft: Have we reached consensus on distributed consensus?
    - Paxos in 25 lines
4. Case study: Google
    - Web search for a planet: The Google Cluster Architecture
    - Building Software Systems at Google and Lessons Learned (video 1:22:44)
5. The Chubby lock service for loosely-coupled distributed systems
6. Megastore: Providing Scalable, Highly Available Storage for Interactive Services
7. Spanner: Google's Globally-Distributed Database
8. Chord: A Scalable Peer-to-peer Lookup Service for Internet Applications
9. 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)
10. Calvin: Fast Distributed Transactions for Partitioned Database Systems
    - Recommended: skim this paper first: The Case for Determinism in Database Systems
11. MapReduce: Simplified Data Processing on Large Clusters
12. Case study: Twitter
    - Real-Time Delivery Architecture at Twitter (56 minute video)
    - Decomposing Twitter: Adventures in Service-Oriented Architecture (50 minute video)
13. Dynamo: Amazon's Highly-available Key-value Store
14. Resilient Distributed Datasets: A Fault-Tolerant Abstration for In-Memory Cluster Computing

---

Presentations

- Conflict-free Replicated Data Types
- Managing Update Conflicts in Bayou, a Weakly Connected Replicated Storage System
- Practical Byzantine Fault Tolerance
- Impossibility of Distributed Consensus with One Faulty Process
- A Note on Distributed Computing
- The Byzantine Generals Problem (Preston)
- Session Guarantees for Weakly Consistent Replicated Data
- CAP Twelve Years Later: How the "Rules" Have Changed (Duy, Minh)
- Distributed Snapshots: Determining Global States of Distributed Systems
- Life beyond Distributed Transactions: an Apostate's Opinion
- Scale and Performance in a Distributed File System (AFS)
- Petal: Distributed Virtual Disks (Travis, Polina)
- On Designing and Deploying Internet-Scale Services
- Dapper, a Large-Scale Distributed Systems Tracing Infrastructure
- PNUTS: Yahoo!'s hosted data serving platform
- Mesa: Geo-Replicated, Near Real-Time, Scalable Data Warehousing
- High-Availability at Massive Scale: Building Google's Data Infrastructure for Ads
- Twitter Heron: Stream Processing at Scale (David, Kaleb)
- Large-scale Incremental Processing Using Distributed Transactions and Notifications
- F1: A Distributed SQL Database That Scales
- Paxos Made Live—An Engineering Perspective
- Flexible Paxos: Quorum intersection revisited
- Large-scale cluster management at Google with Borg (Max, Kendall, Sam)
- Time, Clocks, and the Ordering of Events in a Distributed System
- Exploiting virtual synchrony in distributed systems