CSCI 5817: Database Systems
Fall 2017

Instructor: Dr. Frank W. Miller
Office: ECOT 734
Email: frank.miller@colorado.edu

Office Hours: MW 11:30a-12:30p
Available other times as needed, send me an email or stop by
Class Time/Location: MWF 3:00p-3:50p FLMG 154

Description:
This course covers some advanced topics in Database Systems. The course is primarily focused on advanced concepts associated with implementation and operations. While much of the discussion of these topics is framed with relational databases, a significant portion of the course covers design elements associated with contemporary database designs including key-value stores and graph databases.

DATABASE OPERATIONS
- System Failures
- Concurrency Control
- Transaction Management

DISTRIBUTED DATABASE PROBLEMS
- Distributed Databases and Query Processing
- Distributed Locking and Commit
- Consistent Hashing

NON-RELATIONAL DATABASES
- OLAP vs. OLTP
- Complex Data Types
- Key-Value Stores
- Graph Databases

Most of the operations material will be taken from the textbook. The remainder of the material will be based on academic papers.

Reading:

Textbook: Garcia-Molina, H., Ullman, J. D., and Widom, J., *Database Systems: The Complete Book, 2nd Edition*, Pearson/Prentice-Hall, 2009.

Lamport, L., Shostak, R., and Pease, M., "The Byzantine General's Problem", *ACM Trans. On Programming Languages, 4*, 3 (Jul. 1982).

Lamport L., "Time, Clocks, and the Ordering of Events in a Distributed System". *CACM 21*, 7 (Jul. 1978).

Lamport, L., "A New Solution of Dijkstra's Concurrent Programming Problem", *CACM, 17*, 8 (Aug. 1974).

Ricart, G. and Agrawala A., "An Optimal Algorithm for Mutual Exclusion in Computer Networks", *Communications of the ACM (CACM) 24*, 1 (Jan. 1981)

Stonebraker, M., "SQL Databases v. NoSQL Databases", *CACM 53*, 4 (Apr. 2010).

Karger, D. et. Al., “Consistent Hashing and Random Trees: Distrubuted Caching Protocols for Relieving Hot Spots on the World Wide Web”, *Proceedings of the 29th ACM Symposium on Theory of Computing*, May 1997.

Stoica, I., et. Al., “Chord: A Scalable Peer-to-Peer Lookup Service for Internet Applications”, *Proceedings of ACM SIGCOMM*, 2001.

Chang, F, et. Al., "Bigtable: A Distributed Storage System for Structured Data", *Symposium on Operating Systems Design and Implementation (OSDI)*, USENIX, 2006.

Corbett, J., et. Al., “Spanner: Google’s Globally-Distributed Database”, *Operating Systems Design and Implementation (OSDI)*, 2012.

DeCandia, G., et. Al., "Dynamo: Amazon's Highly Available Key-value Store", *Symposium on Operating System Principles (SOSP)*, ACM, 2007.

Angles, R. and Gutierrez, C., "Survey of Graph Database Models", *ACM Computing Surveys, Vol. 40*, No. 1, Article 1, Feb. 2008.

Graph Databases: New Opportunities for Connected Data, 2nd Ed., Robinson, I., Webber, J., and Eifrem, E., O'Reilly, 2015.

Miller, F. W., "A Graph Model for Database Implementation", unpublished, 2016.

Course Website:
Please enroll ASAP in the Moodle course web page. Nearly all your class interactions will be available through Moodle. The enrollment key is **consistent**

The Moodle site includes two discussion forums. The first is used by the instructors to broadcast to the students. The other is a discussion forum to be used between students. The instructor watches this forum and will occasionally answer questions. The students are encouraged to jump in and answer questions before I get there!

Grading and Policies:

Midterm	20%
Final	30%
Homeworks	30% (probably 4)
Project	20% (big programming project)

Plagiarism policy.

Disability Policy

Religious Observances Policy

Discrimination and Sexual Harassment Policy

Classroom Behavior Policy