CSCI 585 - Database Systems

Units: 4
Fall 2024 - Lecture Friday 2-4:20 pm
Discussion Friday 4:30-5:20 pm
Room: SGM 123

Instructor: Shahram Ghandeharizadeh
Office: SAL 208
Office Hours: Mon 2-3:30 pm, Thur 3-4 pm
Contact Info: shahram@usc.edu, 213-740-4781.

Course Description

This course introduces students to parallel database management systems that scale in a data center and across data centers, physical organization of data for enhanced performance and availability, NoSQL, cache augmented database systems and memory management, disaggregated database systems, transactional storage managers, and data analytics. The discussion sections introduce basics of transactions with ACID properties, a file system, SDM and relational data models, SQL, persistent index structures such as hashing and B+-tree, relational algebra, SQL, normal forms, fundamentals of a magnetic disk drive and solid state drives including their working details, algorithms such as LRU, concurrency control protocols such as locking and time-stamp, and crash recovery protocols such as logging. We review an actor model named FLOW for expedited development of multi-node database systems and study how it is used in FoundationDB.

Guest lecturers include leading database researchers and practitioners. Confirmed presenters include Dr. Haoyu Huang (Google AlloyDB on 9/20), Dr. Hieu Nguyen (eBay NuGraph on 9/27), Dr. Doug Terry (Amazon DynamoDB on 10/4), and Dr. Michael Carey (Couchbase Capella Columnar on 11/15).

Learning Objectives

By the end of the semester, students should have a grasp of fundamental concepts, design decisions, protocols and algorithms to build a database system using off-the-shelf software and hardware components. They should understand analytics and ACID semantics and their application to both storage managers, database systems, and applications. Finally, they should be able to provide an abstraction of today's hardware in support of database systems.

Prerequisite(s): Principles of Software Development (CSCI 201).

Recommended Preparation: A good understanding of in-memory data structures such as records and their organizations in arrays, trees, hash tables, etc. Knowledge of relational databases, SQL, relational algebra and physical database design is required. We will cover the relational data model+algebra, SQL, physical database design in the discussion sections.

Course Notes:

- All lecture material will be posted on the USC Brightspace.
- Grading breakdown
  - Exam 1, Oct 18, 2024: 40% (DEN 50%)
  - Exam 2, Dec 6, 2024: 40% (DEN 50%)
  - Class Participation: 20% (DEN 0%)
- DEN students are exempt from in-class participation. Each Exam counts toward 50% of their grade.
- This course has no comprehensive final exam. Exam 2 is held during the last lecture and covers material presented since Exam 1.

Communication
This course uses the USC Brightspace system to disseminate lectures and assigned readings, and make announcements.

Technological Proficiency and Hardware/Software Required

Required Readings and Supplementary Materials
This course assumes students are Internet savvy and capable of using it as a digital library. This includes use of ACM/IEEE/Springer digital libraries. All USC students have access to these digital collections automatically.

Grading Breakdown
Based on prior experience teaching this course, students who attend lectures do very well. In specific, students who skip lectures do poorly on Exam 1 and spend the second half of the semester concerned with their grade. To prevent this, classroom participation is a requirement for this course. It is worth 20% of the final grade for on campus students. *DEN students are exempt from in-class participation. Each Exam counts toward 50% of their grade.*

Assessment Tool (assignments)	Points	% of Grade
Exam 1, Oct 18, 2024	100	40%
Exam 2, Dec 6, 2024	100	40%
Participation	100	20%
TOTAL		**100%**

Class Participation: How it Works?
At the start of each lecture, you are provided with 1 to 3 questions. Write your answer on a piece of paper and give it to the TA at the end of the lecture. DEN students are exempt with each Exam counting as 50% of their grades.

Grading Scale
The final letter grade is based on a curve.

Course-specific Policies (Assignment Submission, Grading Timeline, Late work, and Technology)
All deadlines are firm. No late submissions or e-mail submissions are accepted for class participation.

Attendance
Based on prior experience, students who do not attend lectures perform poorly on exams. CSCI 585 accommodates student athletes with approved Travel Request Letters and students who give advance notice of religious observation.

Assigned Reading

Week 1: Introductions, Semantic Data Model & Relational Data Model

Assigned reading: Michael Hammer and Dennis McLeod. The Semantic Data Model: A Modelling Mechanism for Data Base Applications. SIGMOD 1978. https://doi.org/10.1145/509252.509264

Discussion

Yazdani, Nima & Alimohammadzadeh, Hamed & Ghandeharizadeh, Shahram. (2023). A Conceptual Model of Intelligent Multimedia Data Rendered using Flying Light Specks. https://www.holodecks.quest/_files/ugd/fb2888_4cf85d8951b241d1a86a378af2144cdb.pdf?index=true (Presented by CSCI 585 TA, Nima Yazdani)

W. Kent. 1983. A Simple Guide to Five Normal Forms in Relational Database Theory. Communications of the ACM, Volume 25, Issue 2, February 1983. https://doi.org/10.1145/358024.358054

D. D. Chamberlin, "Early History of SQL," in IEEE Annals of the History of Computing, vol. 34, no. 4, pp. 78-82, Oct.-Dec. 2012, doi: https://doi.org/10.1109/MAHC.2012.61

Week 2: Parallel Relational Database Management Systems

Assigned Reading: D. J. DeWitt, S. Ghandeharizadeh, D. A. Schneider, A. Bricker, H. -I. Hsiao and R. Rasmussen, "The Gamma database machine project," in IEEE Transactions on Knowledge and Data Engineering, vol. 2, no. 1, pp. 44-62, March 1990, doi: https://doi.org/10.1109/69.50905

Discussion

FLOW (Presented by CSCI 585 TA Hamed Alimohammadzadeh):

Read https://apple.github.io/foundationdb/flow.html,

Look at https://github.com/gyming1849/flow_compile and try to run it.

Transactions and Crash Recovery (Presented by CSCI 585 TA Quankai Gao). Textbook material, e.g., Look at Chapters 15 and 17 of *Database Systems Concepts*, 4th edition, by Silberschatz, Korth, Sudarshan, publisher Morgan Kaufma.

Week 3: Physical Data Design

Assigned Reading: Shahram Ghandeharizadeh and David J. DeWitt. 1990. Hybrid-Range Partitioning Strategy: A New Declustering Strategy for Multiprocessor Database Machines. In Proceedings of the 16th International Conference on Very Large Data Bases (VLDB '90). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 481–492. https://www.vldb.org/conf/1990/P481.PDF

Discussion

Transactions: Concurrency Control Protocols

Week 4: Separation of Storage from Processing, Guest Lecture by Dr. Haoyu Huang of Google

Assigned Reading: Haoyu Huang and Shahram Ghandeharizadeh. 2021. Nova-LSM: A Distributed, Component-based LSM-tree Key-value Store. In Proceedings of the 2021 International Conference on Management of Data (SIGMOD '21). Association for Computing Machinery, New York, NY, USA, 749–763. https://doi.org/10.1145/3448016.3457297

Discussion
Google's AlloyDB

Week 5: Transactional Storage Managers, Guest Lecture by Dr. Hieu Nguyen of eBay

Assigned Reading: Jingyu Zhou, Meng Xu, Alexander Shraer, Bala Namasivayam, Alex Miller, Evan Tschannen, Steve Atherton, Andrew J. Beamon, Rusty Sears, John Leach, Dave Rosenthal, Xin Dong, Will Wilson, Ben Collins, David Scherer, Alec Grieser, Young Liu, Alvin Moore, Bhaskar Muppana, Xiaoge Su, and Vishesh Yadav. 2022. FoundationDB: A Distributed Key Value Store. SIGMOD Rec. 51, 1 (March 2022), 24–31. https://doi.org/10.1145/3542700.3542707

Discussion

Ghandeharizadeh, S., Bernstein, P.A., Borthakur, D., Huang, H., Menon, J., Puri, S. (2023). Disaggregated Database Management Systems. In: Nambiar, R., Poess, M. (eds) Performance Evaluation and Benchmarking. TPCTC 2022. Lecture Notes in Computer Science, vol 13860. Springer, Cham. From the USC secure network, use the following link to access the paper: https://libcatalog.usc.edu/discovery/fulldisplay?context=PC&vid=01USC_INST:01USC&search_scope=MyInst_and_CI&tab=Everything&docid=cdi_scopus_primary_2_s2_0_85152533804

Shahram Ghandeharizadeh and Hieu Nguyen. 2019. Design, Implementation, and Evaluation of Write-Back Policy with Cache Augmented Data Stores. VLDB 2019. https://www.vldb.org/pvldb/vol12/p836-ghandeharizadeh.pdf

Week 6: NoSQL Database Systems, Guest Lecture by Dr. Doug Terry of Amazon

Assigned Reading: Joseph Idziorek and Alex Keyes and Colin Lazier and Somu Perianayagam and Prithvi Ramanathan and James Christopher Sorenson III and Doug Terry and Akshat Vig. 2023. Distributed Transactions at Scale in Amazon DynamoDB. 2023 USENIX Annual Technical Conference (USENIX ATC 23). https://www.usenix.org/conference/atc23/presentation/idziorek

Discussion

Mostafa Elhemali and Niall Gallagher and Nick Gordon and Joseph Idziorek and Richard Krog and Colin Lazier and Erben Mo and Akhilesh Mritunjai and Somasundaram Perianayagam and Tim Rath and Swami Sivasubramanian and James Christopher Sorenson III and Sroaj Sosothikul and Doug Terry and Akshat Vig. 2022. Amazon DynamoDB: A Scalable, Predictably Performant, and Fully Managed NoSQL Database Service, 2022 USENIX Annual Technical Conference (USENIX ATC 22). https://www.usenix.org/conference/atc22/presentation/elhemali

Jingyu Zhou, Meng Xu, Alexander Shraer, Bala Namasivayam, Alex Miller, Evan Tschannen, Steve Atherton, Andrew J. Beamon, Rusty Sears, John Leach, Dave Rosenthal, Xin Dong, Will Wilson, Ben Collins, David Scherer, Alec Grieser, Young Liu, Alvin Moore, Bhaskar Muppana, Xiaoge Su, Vishesh Yadav: FoundationDB: A Distributed Unbundled Transactional Key Value Store. SIGMOD Conference 2021: 2653-2666 https://www.foundationdb.org/files/fdb-paper.pdf

Week 7: Fall Recess

Week 8: Exam 1
Exam 1 covers Week 1-6 assigned readings.

Week 9: Benchmarks
Assigned Reading: S. Barahmand and S. Ghandeharizadeh. BG: A Benchmark to Evaluate Interactive Social Networking Actions. Proceedings of the biennial Conference on Innovative Data Systems Research (CIDR), Asilomar, California, January 6-9, 2013.
https://www.cidrdb.org/cidr2013/Papers/CIDR13_Paper93.pdf

Discussion
Brian F. Cooper, Adam Silberstein, Erwin Tam, Raghu Ramakrishnan, and Russell Sears. 2010. Benchmarking cloud serving systems with YCSB. In Proceedings of the 1st ACM symposium on Cloud computing (SoCC '10). Association for Computing Machinery, New York, NY, USA, 143-154. https://doi.org/10.1145/1807128.1807152

Week 10: Memory Management
Assigned Reading: Shahram Ghandeharizadeh, Sandy Irani, Jenny Lam, and Jason Yap. 2014. CAMP: A Cost Adaptive Multi-Queue Eviction Policy for Key-Value Stores. In Proceedings of the 15th International Middleware Conference (Middleware '14). Association for Computing Machinery, New York, NY, USA, 289–300.
https://doi.org/10.1145/2663165.2663317

Discussion
Elizabeth J. O'Neil, Patrick E. O'Neil, and Gerhard Weikum. 1993. The LRU-K page replacement algorithm for database disk buffering. SIGMOD Rec. 22, 2 (June 1, 1993), 297-306. https://doi.org/10.1145/170036.170081

Week 11: Cache Augmented Database Management Systems
Assigned Reading: Shahram Ghandeharizadeh, Jason Yap, and Hieu Nguyen. 2014. Strong consistency in cache augmented SQL systems. Proceedings of the 15th International Middleware Conference (Middleware '14). Association for Computing Machinery, New York, NY, USA, 181–192.
https://doi.org/10.1145/2663165.2663318 OR using USC DL
https://dl-acm-org.libproxy1.usc.edu/doi/pdf/10.1145/2663165.2663318

Discussion
Shahram Ghandeharizadeh, Sandy Irani, and Jenny Lam. 2015. Cache replacement with memory allocation. In Proceedings of the Meeting on Algorithm Engineering & Experiments (ALENEX '15). Society for Industrial and Applied Mathematics, USA, 1-9. https://epubs.siam.org/doi/10.1137/1.9781611973754.1

Week 12: Data Analytics, Guest Lecture by Dr. Michael Carey of Couchbase

Assigned Reading: Michael J. Carey, Don Chamberlin, Almann Goo, Kian Win Ong, Yannis Papakonstantinou, Chris Suver, Sitaram Vemulapalli, Till Westmann. SQL++: We Can Finally Relax! ICDE 2024: 5501-5510.
https://doi.org/10.1109/ICDE60146.2024.00438

Discussion

William Spoth, Bahareh Sadat Arab, Eric S. Chan, Dieter Gawlick, Adel Ghoneimy, Boris Glavic, Beda Christoph Hammerschmidt, Oliver Kennedy, Seokki Lee, Zhen Hua Liu, Xing Niu, Ying Yang. Adaptive Schema Databases. CIDR 2017.
https://willspoth.github.io/pdf/Will_Spoth_Adaptive_Databases.pdf

Week 13: MapReduce

Assigned Reading: Jeffrey Dean and Sanjay Ghemawat. MapReduce: Simplified Data Processing on Large Clusters. Sixth Symposium on Operating System Design and Implementation (OSDI)2004.
https://dl.acm.org/doi/pdf/10.1145/1327452.1327492

Discussion

Juwei Shi, Yunjie Qiu, Umar Farooq Minhas, Limei Jiao, Chen Wang, Berthold Reinwald, and Fatma Özcan. 2015. Clash of the titans: MapReduce vs. Spark for large scale data analytics. Proc. VLDB Endow. 8, 13 (September 2015), 2110-2121.
https://doi.org/10.14778/2831360.2831365

Week 14: Thanksgiving Break

Week 15: Exam 2

Exam 2 covers Week 9-13 assigned readings.

Fall 2024 Course Schedule: A Weekly Breakdown

	Topics/Daily Activities	**Readings/Preparation**	**Key Concepts**
Week 1 8/30	Introduction, SDM, Relational	SDM, Conceptual model. Discussion led by Nima Yazdani.	Conceptual data models
Week 2 9/6	Parallel DBMSs	Gamma. Actor/Flow discussion led by Yiming Gao.	Parallel query processing, crash recovery, shared-nothing architecture
Week 3 9/13	Data Declustering	Hybrid Range Partitioning Strategy	Sharding, Load Balancing

Week 4 9/20	Disaggregated DBMSs	Nova-LSM, Google's AlloydB. Guest lecture by Dr. Haoyu Huang, Google.	Separation of storage from processing, shared disk architecture.
Week 5 9/27	Xact Storage Managers	FoundationDB, Guest Lecture by Dr. Hieu Nguyen, eBay.	OCC, MVCC, Fat Clients, DBMS as a collection of microservices
Week 6 10/4	NoSQL + Xacts at Scale	NoSQL DBMSs. Guest lecture by Dr. Doug Terry, Amazon.	NoSQL
Week 7 10/11	Fall Recess	Fall Recess	Fall Recess
Week 8 10/18	**Review for Exam 1 (1 hour) + 1 hour Exam 1 is on the same day**	Exam 1	**Review for Exam 1 + Exam 1.**
Week 9 10/25	Benchmarks	BG Social Networking Benchmark	Performance metrics, Closed and Open Simulation Models
Week 10 11/1	Memory Mgmt	CAMP and LRU	Low complexity memory mgmt algorithms that consider size and cost.
Week 11 11/8	Cache Augmented Database Mgmt Systems	IQ Framework	Impact of caches on transactions
Week 12 11/15	Data Analytics	Columnar stores. Guest Lecture by Dr. Michael Carey, Couchbase.	In-the-cloud JSON data analytics.
Week 13 11/22	MapReduce	Programming paradigm for parallelism.	Divide-and-Conquer, Resilient to failures.
Week 14 11/29	Thanksgiving Break	Thanksgiving Break	Thanksgiving Break
Week 15 12/6	**Review for Exam 2 & Exam 2**		**Exam 2** 12/6

Statement on Academic Conduct and Support Systems

Academic Conduct:

Plagiarism – presenting someone else's ideas as your own, either verbatim or recast in your own words – is a serious academic offense with serious consequences. Please familiarize yourself with the discussion of plagiarism in SCampus in Part B, Section 11, "Behavior Violating University Standards" policy.usc.edu/scampus-part-b. Other forms of academic dishonesty are equally unacceptable. See additional information in SCampus and university policies on Research and Scholarship Misconduct.

Students and Disability Accommodations:

USC welcomes students with disabilities into all of the University's educational programs. The Office of Student Accessibility Services (OSAS) is responsible for the determination of appropriate accommodations for students who encounter disability-related barriers. Once a student has completed the OSAS process (registration, initial appointment, and submitted documentation) and accommodations are determined to be reasonable and appropriate, a Letter of Accommodation (LOA) will be available to generate for each course. The LOA must be given to each course instructor by the student and followed up with a discussion. This should be done as early in the semester as possible as accommodations are not retroactive. More information can be found at osas.usc.edu. You may contact OSAS at (213) 740-0776 or via email at osasfrontdesk@usc.edu.

Support Systems:

Counseling and Mental Health - (213) 740-9355 – 24/7 on call
studenthealth.usc.edu/counseling
Free and confidential mental health treatment for students, including short-term psychotherapy, group counseling, stress fitness workshops, and crisis intervention.

National Suicide Prevention Lifeline - 1 (800) 273-8255 – 24/7 on call
suicidepreventionlifeline.org
Free and confidential emotional support to people in suicidal crisis or emotional distress 24 hours a day, 7 days a week.

Relationship and Sexual Violence Prevention Services (RSVP) - (213) 740-9355(WELL), press "0" after hours – 24/7 on call
studenthealth.usc.edu/sexual-assault
Free and confidential therapy services, workshops, and training for situations related to gender-based harm.

Office for Equity, Equal Opportunity, and Title IX (EEO-TIX) - (213) 740-5086
eeotix.usc.edu
Information about how to get help or help someone affected by harassment or discrimination, rights of protected classes, reporting options, and additional resources for students, faculty, staff, visitors, and applicants.

Reporting Incidents of Bias or Harassment - (213) 740-5086 or (213) 821-8298
usc-advocate.symplicity.com/care_report
Avenue to report incidents of bias, hate crimes, and microaggressions to the Office for Equity, Equal Opportunity, and Title for appropriate investigation, supportive measures, and response.

The Office of Student Accessibility Services (OSAS) - (213) 740-0776
osas.usc.edu
OSAS ensures equal access for students with disabilities through providing academic accommodations and auxiliary aids in accordance with federal laws and university policy.
USC Campus Support and Intervention - (213) 821-4710
campussupport.usc.edu
Assists students and families in resolving complex personal, financial, and academic issues adversely affecting their success as a student.

Diversity, Equity and Inclusion - (213) 740-2101
diversity.usc.edu

Information on events, programs and training, the Provost's Diversity and Inclusion Council, Diversity Liaisons for each academic school, chronology, participation, and various resources for students.

USC Emergency - UPC: (213) 740-4321, HSC: (323) 442-1000 – 24/7 on call
dps.usc.edu, emergency.usc.edu
Emergency assistance and avenue to report a crime. Latest updates regarding safety, including ways in which instruction will be continued if an officially declared emergency makes travel to campus infeasible.

USC Department of Public Safety - UPC: (213) 740-6000, HSC: (323) 442-120 – 24/7 on call
dps.usc.edu
Non-emergency assistance or information.

Office of the Ombuds - (213) 821-9556 (UPC) / (323-442-0382 (HSC)
ombuds.usc.edu
A safe and confidential place to share your USC-related issues with a University Ombuds who will work with you to explore options or paths to manage your concern.

Occupational Therapy Faculty Practice - (323) 442-3340 or otfp@med.usc.edu
chan.usc.edu/otfp
Confidential Lifestyle Redesign services for USC students to support health promoting habits and routines that enhance quality of life and academic performance.