Intro to Database Systems (15-445/645)

26 Final Review

Carnegie Mellon University

FALL 2022

Andy Pavlo

ADMINISTRIVIA

Project #4 is due **Sunday Dec 11th @ 11:59pm**

→ Extra Office Hours: **Saturday Dec 10th @ 3:00-5:00pm**

→ Location: GHC 5209, GHC 5211

Final Exam is **Friday Dec 16th @ 1:00pm.**

→ Study guide will be posted next week.

→ Location: GHC 4401

SPRING 2023

Charlie (aka "Cutty") is recruiting impressionable TAs for 15-445/645 in Spring 2023.

→ All BusTub projects will remain in C++.

https://forms.gle/AvjfUtSaWtrNiJMXA

COURSE EVALS

Your feedback is strongly needed:
→ https://cmu.smartevals.com
→ https://www.ugrad.cs.cmu.edu/ta/F22/feedback/

Things that we want feedback on:
→ Homework Assignments
→ Projects
→ Reading Materials
→ Lectures

OFFICE HOURS

Andy's hours:

→ Tuesday Dec 13th @ 12:30-1:30pm (GHC 9019)
→ Thursday Dec 15th @ 8:00-9:00pm (Zoom)
→ https://savvycal.com/pavlo/f22-445
→ Or by appointment

All TAs will have their regular office hours up to and including Saturday Dec 10th

FINAL EXAM

Who: You

What: Final Exam

Where: GHC 4401

When: Thursday Dec 16th @ 1:00pm

Why: https://youtu.be/8tuoIO4CxOw

Email Andy if you need special accommodations.

https://15445.courses.cs.cmu.edu/fall2022/final-guide.html

FINAL EXAM

What to bring:

→ CMU ID
→ Pencil + Eraser (!!!)
→ Calculator (cellphone is okay)
→ One 8.5x11" page of handwritten notes (double-sided)

What not to bring:

→ NFT-themed Clothing

STUFF BEFORE MID-TERM

SQL

Buffer Pool Management

Hash Tables

B+Trees

Storage Models

Inter-Query Parallelism

QUERY OPTIMIZATION

Heuristics
→ Predicate Pushdown
→ Projection Pushdown
→ Nested Sub-Queries: Rewrite and Decompose

Statistics
→ Cardinality Estimation
→ Histograms

Cost-based search

TRANSACTIONS

ACID

Conflict Serializability:
→ How to check?
→ How to ensure?

View Serializability

Recoverable Schedules

Isolation Levels / Anomalies

TRANSACTIONS

Two-Phase Locking
→ Rigorous vs. Non-Rigorous
→ Deadlock Detection & Prevention

Multiple Granularity Locking
→ Intention Locks

TRANSACTIONS

Timestamp Ordering Concurrency Control
→ Thomas Write Rule

Optimistic Concurrency Control
→ Read Phase
→ Validation Phase
→ Write Phase

Multi-Version Concurrency Control
→ Version Storage / Ordering
→ Garbage Collection

CRASH RECOVERY

Buffer Pool Policies:
→ STEAL vs. NO-STEAL
→ FORCE vs. NO-FORCE

Write-Ahead Logging

Logging Schemes

Checkpoints

ARIES Recovery
→ Log Sequence Numbers
→ CLRs

DISTRIBUTED DATABASES

System Architectures

Replication

Partitioning Schemes

Two-Phase Commit

TOPICS NOT ON EXAM!

Snowflake

Embedded Logic (Lecture #24)

Details of specific database systems (e.g., Postgres)