

15-721

DATABASE SYSTEMS



Lecture #24 – End of Semester

Andy Pavlo // Carnegie Mellon University // Spring 2016

TODAY'S AGENDA

End-of-Semester Logistics

Final Review Topics

MemSQL Tech Talk

ADMINISTRIVIA

Final Exam: April 27th @ 12:00pm

→ Three short-essay questions.

“Final” Presentations: May 6th @ 1:00pm

→ 10 minutes per group

→ Food and prizes for everyone!

Code Reviews:

→ Pull Request Date: **May 6, 2016 @ 11:59pm**

→ Review Date: **May 9, 2016 @ 11:59pm**

PROJECT #3 – CODE DROP

Your code must merge into the master branch without any conflicts.

You must address all comments from your code review.

The project includes test cases that correctly verify that implementation is correct.

Implementation documentation in a separate Markdown file.

Due Date: May 11, 2016

EXTRA CREDIT

Website is still not ready.

I will provide a JSON schema that you can use to write your article.

Due Date: May 11, 2016 @ 11:59pm

COURSE EVALUATION

You guys were guinea pigs.

Please be sure to fill out the course evaluation to help me improve the course for next year.

Things that I already know:

- Project #2 was hard.
- Autolab was a pain in the ass.
- I need to provide better grades/feedback.
- My personal hygiene has been described as offensive.

FINAL EXAM

Three questions.

Short essays.

Closed notes.

TOPICS

Multi-version Concurrency Control

Storage Models

Query Processing Models

→ Tuple-at-a-time (Volcano)

→ Materialized (MonetDB)

→ Vectorized (VectorWise)

Join Algorithms

Logging Schemes

OLTP + OLAP Indexes

TOPICS

Compression Schemes

Execution Optimizations

- Query Compilation / Code Generation
- Scan Sharing
- Vectorization

Hardware Optimizations

- Cache Locality
- Vectorization / SIMD