



# **15-721**

## **ADVANCED DATABASE SYSTEMS**

**Lecture #25 – End of Semester +  
Impala/Kudu Tech Talk**

@Andy\_Pavlo // Carnegie Mellon University // Spring 2017

# TODAY'S AGENDA

---

End-of-Semester Logistics

Final Review Topics

Impala Tech Talk



# ADMINISTRIVIA

---

## **Final Exam: May 4<sup>th</sup> @ 12:00pm**

→ Multiple choice + short-answer questions.

## **Code Review #2: May 4<sup>th</sup> @ 11:59pm**

→ We will use the same group pairings as before.

## **Final Presentations: May 9<sup>th</sup> @ 5:30pm**

→ **WEH Hall 7500**

→ 12 minutes per group

→ Food and prizes for everyone!



# 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 15, 2016**

# EXTRA CREDIT

---

Website got broken this weekend.

All data is preserved.

If we can't fix it, I will provide a JSON schema that you can use to write your article.

Due Date: **May 15, 2016 @ 11:59pm**



# COURSE EVALUATION

---

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

Things that I already know:

- Autolab was a pain in the ass.
- Make it easier to program on a Mac.
- Source code conflicts were annoying.
- My personal hygiene has been described as offensive.



# FINAL EXAM

---

Three questions.

Closed notes.



# TOPICS

---

OCC + MVCC

Storage Models

Query Processing Models

→ Tuple-at-a-time (Volcano)

→ Materialized (MonetDB)

→ Vectorized (Vectorwise)

Join Algorithms

Logging Schemes

OLTP + OLAP Indexes





# TOPICS

---

Optimizer Implementation

Compression Schemes

Execution Optimizations

→ Query Compilation / Code Generation

→ Vectorization

Hardware Optimizations

→ Cache Locality

→ Vectorization / SIMD

