


Practice Problems and Wrap Up

Introduction to Databases
CompSci 316 Spring 2017



Announcements (Wed., Apr. 26)

- **Project code** due tonight: Wednesday, April 26, 11:59 pm
 - See all announcements about project report and demo on piazza
 - **Please do a final coordination as a group**
- **Final Exam**
 - May 2 (next Tuesday), 2-5 pm, in class
 - Everything covered in the class (up to last lecture) is included
 - **If you need special accommodation, please email me**
 - **Open book, open notes**
 - any written material is allowed
 - no electronic gadget (phone, laptop, calculator) is allowed
- **Practice final exam uploaded**
 - under sakai -> resources -> practiceproblems
 - prepare and try it first in 3 hours before seeing solutions!
- **Tomorrow's demonstrations**
 - Please plan for strict 15-16 mins of demonstration
 - leave 3-4 mins for questions

Today

- 10 Practice problems and wrap up
- A few mins for you to try each problem before we discuss solutions
- **Disclaimer: Numerous other possibility – this is not a representative of what may or may not appear**
- 10 mins break for filling out course evaluations after 5 problems
- If we have time, we can review any topic you want

Practice Problems are on Sakai

- Resources -> PracticeProblems -> Lecture27Probs

What did we learn in this class?

- Relational Model, Query Languages, and Database Design
 - SQL
 - RA
 - E/R diagram
 - Normalization
- DBMS Internals
 - Storage
 - Indexing
 - Query Evaluation
 - External sort
 - Join Algorithms
 - Query Optimization

- Transactions
 - Basic concepts
 - Concurrency control
 - Recovery
- XML
 - DTD and XML schema
 - XPath and XQuery
 - Relational Mapping
- Selected other topics
 - Parallel DBMS
 - Map Reduce
 - Data Mining
 - Data Warehousing
 - Distributed DBMS
 - NOSQL and Column Store

Many other research areas : 1/2

- Data analytics, interactive exploration, visualization
- Big data systems (Spark, Map Reduce, ...)
- Approximate Query processing
- Storage and new hardware (GPU, FPGA, memory..)
- Social networks
- Data streams
- Incremental View Maintenance
- Tree and Graph Databases
- Data integration
- Data cleaning

Many other research areas : 2/2

- Query optimization
- NOSQL and NewSQL
- User preferences
- Machine learning in databases
- Spatial and temporal data
- Data privacy and security
- Crowd sourcing
- Data provenance
- Inconsistent and incomplete databases
- Database theory (logic, algorithms, complexity)
-

More resources

- Check out top database conferences (huge participation from industry and academia)
 - SIGMOD and PODS <http://sigmod2017.org>
 - VLDB <http://www.vldb.org/2017>
 - ICDE <http://icde2017.sdsc.edu>
 - ICDT <http://edbticdt2017.unive.it> (PODS and ICDT are db theory conferences)
- Advanced (graduate) database courses at Duke CS
- Duke CS database group:
https://sites.duke.edu/duke_dbgroup/
- Always happy to talk if you are interested in database related research!

- Please fill out course evaluations!
- Good luck with all your exams!
- and **Thank You!**