

CSE 344: Topics for Final



Pre-Midterm Material

- SQL
 - FWGHOS
 - Nesting
- Datalog
- NoSQL/SQL++
- Relational Algebra

Post-Midterm Material

- Physical Plans/Cost Estimation/Indexing
 - Join/Selection Algorithms
 - Clustered/Unclustered Indexes
- Parallel DB/MapReduce
- Design
 - Table Creation
 - Design Theory
 - Weak Entities
- Transactions
 - Serializable/Conflict Serializable
 - Locking
 - Guarantees

Physical Plans/Cost Estimation/Indexing

- Single Node Joins
 - Nested Loop Joins (Tuple, Block, Block-Nested)
 - Hash Join
 - Sort-Merge Join
- Selectivity Factor
- Indexed Selection and Join
 - Clustered
 - Unclustered
- General Hardware Semantics
 - Disk
 - Main Memory

Parallel DB/MapReduce

- Shared-Nothing Architecture
- Intra-Operator Parallelism
- Shuffling and Broadcasting for Grouping and Joining
- MapReduce Framework
 - Map tasks
 - Reduce tasks
- Recovery in MapReduce and Spark

Design

- Table Creation
- Relationships
 - 1-1
 - N-1
 - N-N
- Subclassing
- Weak Entities
- Design Theory
 - Functional Dependencies
 - Closures
 - BCNF

Transactions

- **Serializable/Conflict Serializable**
 - Precedence Graphs
 - Conflicts (RW, WR, WW)
- **Locking**
 - 2PL/Strict 2PL
 - Deadlock
- **Guarantees**
 - READ UNCOMMITTED
 - READ COMMITTED
 - REPEATABLE READ
 - SERIALIZABLE
 - Phantom Problem

General Advice

Make a note sheet! It's a good way to study.

Lots of practice materials linked in last quarter's website.