

CSE 344

JUNE 1ST – FINAL REVIEW

EXAMINATION STATIONS

- **Exam Wednesday June 6**
 - 8:30-10:20
- **One sheet of notes, front and back**
- **Practice exam/solutions out after class**
- **Good luck!**

EXAM LENGTH

- **1:50 minutes**
- **Roughly 50% longer than midterm**
- **Will focus on material from second half of the course**
 - Practice Exam
 - Midterm exam

FIRST HALF TOPICS

- **Databases**
 - Motivations and definitions
- **Relational Databases**
 - SQL
 - Relational Algebra
 - Datalog
- **Semi-structured Data**
 - Motivations and definitions

FIRST HALF TOPICS

- **Internals**
 - Indexes
 - Physical plans/Cost Estimation
 - Disk I/o
- **Parallel**
 - Shared Nothing
 - Map Reduce

NEW MATERIAL

- **Database Design**
 - E/R Diagrams
 - Constraints
 - 1NF, BCNF
 - Decomposition
 - Loss/Chase test

NEW MATERIAL

- **Transactions**
 - Schedules
 - Concurrent behavior
 - ACID
 - Locking

DATABASE DESIGN

- **E/R Diagrams**

DATABASE DESIGN

- **E/R Diagrams**

- Entities, Relations, Attributes
- Many-to-many, etc
- Weak entities
- Constraints
- Converting to schema

DATABASE DESIGN

- **Normalization**

DATABASE DESIGN

- **Normalization**
 - Functional Dependencies
 - Closures
 - Keys/super keys
 - Normal Forms
 - 1NF, BCNF

DATABASE DESIGN

- **BCNF**

DATABASE DESIGN

- **BCNF**
 - Recognizing BCNF
 - Decomposing
 - Generating schemas from FDs
 - Lossless decomposition
 - Chase test
 - Views (Materialized and Virtual)

TRANSACTIONS

- **Motivations**

TRANSACTIONS

- **Motivations**
 - Concurrent behavior
 - Throughput v. Safety
 - Conflicts
 - ACID
 - Deadlock

TRANSACTIONS

- **Schedules**

TRANSACTIONS

- **Schedules**
 - Serial
 - Serializable
 - Conflict-serializable
 - Precedence Graphs
 - (Recoverable)

TRANSACTIONS

- **Schedulers**

TRANSACTIONS

- **Schedulers**
 - Locking
 - SQLite
 - 2PL, Strict 2PL
 - Granularity (DB v. Records)
 - Phantom problem
 - Isolation level

TRANSACTIONS

- **Design considerations**

TRANSACTIONS

- **Design considerations**
 - Throughput
 - Thrashing
 - Deadlock
 - ACID
 - Transactional v. Analytical

ON THE EXAM

- **Short answer**
 - General/1st half
 - Database Design
 - Transactions
 - *May have differing point values*

LONG ANSWER

- **BCNF Decomposition**
- **E/R Diagrams**
 - Schemas and Diagrams
- **Schedules**
 - Conflict-Serializability

SOME OF THE FOLLOWING

- **From the midterm**
 - SQL/Datalog/RA
 - Cost Estimation
- **Missing from Midterm**
 - SQL++
 - MapReduce (pseudocode only)

SOME OF THE FOLLOWING

- **Database Design**
 - Views/Indexes
 - Workload
 - Isolation levels
 - Client-specification/Design Pipeline

SOME OF THE FOLLOWING

- **Transactions**
 - Producing schedules
 - Verifying schedules
 - Locks and unlocks
 - Constraints

ADVICE

- **Look through the exam first**
 - Try and do easiest questions first
 - Short answer questions *may not* be worth equal amounts
 - Points per on exam
- **Be ready for tricky questions**
 - Most problems in familiar format

FINAL ASSORTED MINUTIAE

- **Office hours (CSE 214)**
 - Monday 12-2
 - Tuesday 2-4
 - These are the only office hours
- **Course evaluations**
 - In by Sunday
- **HW8 Due Tonight 11:30**
 - Up to two late days
 - Please only tag once