Final Exam Review/Topics

ECS 165A Winter 2011

TJ Green

Scope of Exam

- Cumulative, will cover material from entire quarter
- Wednesday 10:30am-12:30pm (2hrs), here
- Similar in flavor to midterm + sample finals (but hopefully less time pressure than midterm)

What *Won't* Be on Final

- SQL assertions
- Details of trigger syntax
  - But there may be high-level/conceptual questions
- Dependency theory, BCNF and 3NF decompositions
- Database connection libraries and PHP
- Stratified negation (Datalog)
  - But Datalog *will* be on final

What *Will* Be on Final

- Database design and the E/R model
- Relational model and relational algebra
- SQL
- Integrity constraints
- Storage and file organization
- Indexing
- Query processing
- Transactions
- Datalog and recursive queries

Database Design and the ER Model

- Basic concepts: entities and entity types, attributes and keys, relationships and relationship types
- Entity-Relationship schema (aka ER diagram)
- Constraints on relationship types
- Design choices
- Enhanced Entity-Relationship model features
- Steps in designing an ER schema
- Translation of an ER schema to tables

Relational Model and Relational Algebra

- Fundamental concepts of the relational model
- Integrity constraints
- Translation of ER diagram -> relational schema
- Relational algebra
- Modifications of the database
  - you *don't* have to know the notation presented at the end of slides for Section 3

SQL

- Basic queries in SQL (select-from-where)
- Set operations on relations
- Nested subqueries
- ~~Null values~~
- Aggregate functions and grouping
- DDL (data definition language)
- Insert, update, and delete statements
- View definitions

Integrity Constraints

- Referential integrity
- "Granularity" of constraint checking
  - per-table versus per-row
- Assertions
- Event-condition-action triggers
  - at the level of pseudocode, don't need to memorize actual syntax

Storage and File Structures

- Tradeoffs of physical storage media
- Magnetic disks, tertiary storage
- Buffer management
- Storage access
- File organization
  - fixed size versus variable length records

Indexing

- Single-level ordered indexes
- Multi-level indexes
- Dense, sparse, primary, secondary
- **B+ tree indexes**
- Index definition in SQL
  - should know the basic syntax

Query Processing

- Catalog information for cost estimation
- Measures of query cost
- Physical operator implementations
  - Selection
  - Join
  - Set operations
- Evaluation and transformation of expressions
- Holistic view of “the life of a query”

Transactions

- ACID requirements
- Violations of isolation: lost updates, dirty reads, inconsistent analysis
- Serializability
- Two-phase locking protocol

Datalog

- Syntax (head, body, subgoal, rule, ...) and semantics
- Safety
- EDBs versus IDBs
- Recursion
- Translating among Datalog, SQL, RA, English