Introduction to CS 4604

Zaki Malik

August 26, 2007

Course Information

- **Instructor**
  Zaki Malik, 2205 CRC KWII, 231-8573, zaki@cs.vt.edu
  - *Office Hours*: 2pm-4pm Mondays and 2pm-3pm Wednesdays

- **Teaching Assistant**
  Haiyan Cheng, McBryde 106, hcheng04@vt.edu
  - Office Hours: 3-5pm Tuesdays and 4-5pm Thursdays

- **Class Meeting Time**
  Tuesdays and Thursdays 5–6:15pm, McBryde 218

- **Keeping in Touch**
  Course web site http://courses.cs.vt.edu/~cs4604
  updated regularly through the semester
  - *Listserv*: cs4604_91770@listserv.vt.edu

Textbook

- **Required**

  A First Course in Database Systems, Ullman and Widom, Prentice Hall. (3rd Ed).

  Web page for the book

  http://www-db.stanford.edu/~ullman/fcdb.html

Course Grading

Homework	30%	5–6
Midterm exam	15%	October 16
Final exam	25%	December 16
Course project	30%	7 assignments

- Project is spread over 7 deliverables
- Projects and homework assignments alternate
- Submit hard copies of homeworks and project assignments at the start of class on the due date
- Each class has required reading. Please consult the course web page
- No Pop-Quizzes ☺

Course Project

- Project overview

  http://courses.cs.vt.edu/~cs4604/Fall08/project/project.html
- 2, or 3 persons per project.
- Project runs the entire semester with regular assignments and a final implementation assignment.
- You are free to suggest a project. The project should not be “overly simple”.
- Send email to Haiyan by 5pm Monday, Sep 01, 2008 stating which project you want to work on.

Why Study Databases?

- Academic
  - Databases involve many aspects of computer science
  - Fertile area of research
  - Three Turing awards in databases
- Programmer
  - a plethora of applications involve using and accessing databases
- Businessman
  - Everybody needs databases => lots of money to be made
- Student
  - Get those last three credits and I don't have to come back to Blacksburg ever again!!!
  - Google, Oracle, Microsoft, etc. will hire me!!
  - Databases sound cool!
  - ???

What Will You Learn in CS4604?

- Implementation
  - How do you build a system such as ORACLE or MySQL?
- Design
  - How do you model your data and structure your information in a database?
- Programming
  - How do you use the capabilities of a DBMS?
- CS 4604 achieves a balance between
  - a firm theoretical foundation to designing moderate-sized databases
  - creating, querying, and implementing realistic databases and connecting them to applications

Course Goals and Outcomes

- Take an English language description and convert it into a working database application.
- Create E/R models from application descriptions.
- Convert E/R models into relational designs.
- Identify redundancies in designs and remove them using normalization techniques.
- Create databases in an RDBMS and enforce data integrity constraints using SQL.
- Write sophisticated database queries using SQL.
- Understand tradeoffs between different ways of phrasing the same query.
- Implement a web interface to a database.

Course Outline

- Weeks 1–5, 13: Query/Manipulation Languages
  - Relational Algebra
  - Data definition
  - Programming with SQL
- Weeks 6–8: Data Modeling
  - Entity-Relationship (E/R) approach
  - Specifying Constraints
  - Good E/R design
- Weeks 9–13: Relational Design
  - The relational model
  - Converting ER to “R”
  - Normalization to avoid redundancy
- Week 14–15: Students’ choice
  - Practice Problems
  - XML
  - Query optimization
  - Data mining

What is a DBMS?

- Database Management System (DBMS) = data + set of instructions to access/manipulate data

- Features of a DBMS
  - Support massive amounts of data
  - Persistent storage
  - Efficient and convenient access
  - Secure, concurrent, and atomic access

- Examples?
  - Search engines, banking systems, airline reservations, corporate records, payrolls, sales inventories.
  - New applications: Wikis, biological/multimedia/scientific/geographic data, heterogeneous data.

Features of a DBMS

- Support massive amounts of data
  - Giga/tera/petabytes
  - Far too big for main memory
- Persistent storage
  - Programs update, query, manipulate data.
  - Data continues to live long after program finishes.
- Efficient and convenient access
  - Efficient: do not search entire database to answer a query.
  - Convenient: allow users to query the data as easily as possible.
- Secure, concurrent, and atomic access
  - Allow multiple users to access database simultaneously.
  - Allow a user access to only to authorized data.
  - Provide some guarantee of reliability against system failures.

A Brief History of DBMS

- The earliest databases (1960s) evolved from file systems
  - File systems
    - Allow storage of large amounts of data over a long period of time
    - File systems do not support:
      - Efficient access of data items whose location in a particular file is not known
      - Logical structure of data is limited to creation of directory structures
      - Concurrent access: Multiple users modifying a single file generate non-uniform results
    - Navigational and hierarchical
    - User programmed the queries by walking from node to node in the DBMS.

- Relational DBMS (1970s to now)
  - View database in terms of relations or tables
  - High-level query and definition languages such as SQL
  - Allow user to specify what (s)he wants, not how to get what (s)he wants

- Object-oriented DBMS (1980s)
  - Inspired by object-oriented languages
  - Object-relational DBMS

The DBMS Industry

- A DBMS is a software system.
- Major DBMS vendors: Oracle, Microsoft, IBM, Sybase
- Free/Open-source DBMS: MySQL, PostgreSQL, Firebird.
  - Used by companies such as Google, Yahoo, Lycos, BASF.
- All are “relational” (or “object-relational”) DBMS.

Example Scenario

- RDBMS = “Relational”DBMS
- The relational model uses relations or tables to structure data
- ClassList relation:

Student	Course	Grade
Hermione Grainger	Potions	A-
Draco Malfoy	Potions	B
Harry Potter	Potions	A
Ron Weasley	Potions	C

- Relation separates the logical view (externals) from the physical view (internals)
- Simple query languages (SQL) for accessing/modifying data
  - Find all students whose grades are better than B.
  - SELECT Student FROM ClassList WHERE Grade >“B”

DBMS Architecture

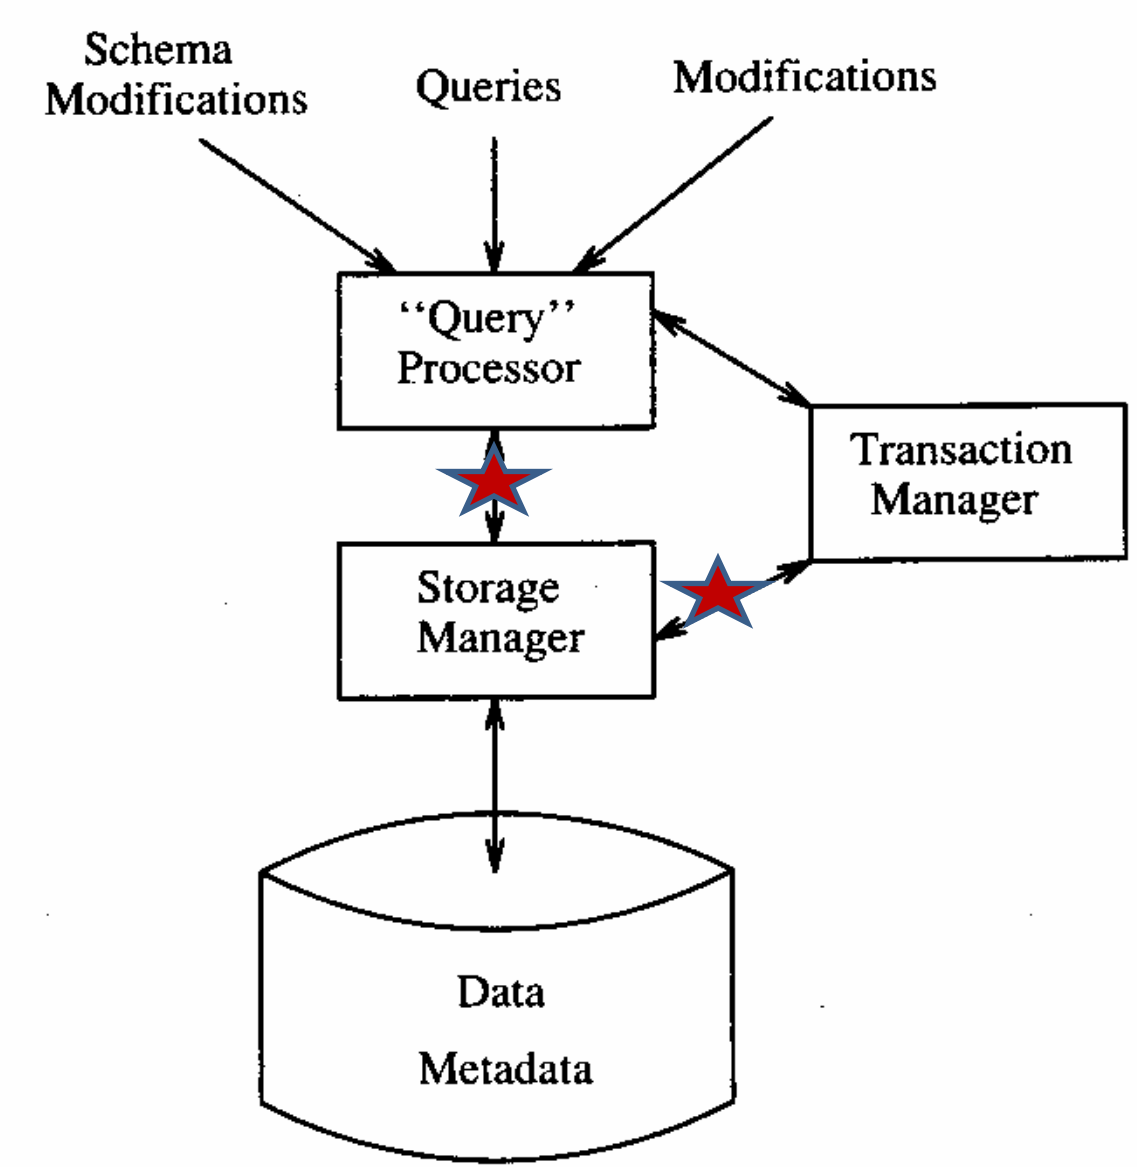

Transaction Processing

- One or more database operations are grouped into a "transaction"

- Transactions should meet the "ACID test"
  - **A**tomicity: All-or-nothing execution of transactions.
  - **C**onsistency: Databases have consistency rules (e.g. what data is valid). A transaction should NOT violate the database's consistency. If it does, it needs to be *rolled back.*
  - **I**solation: Each transaction must appear to be executed as if no other transaction is executing at the same time.
  - **D**urability: Any change a transaction makes to the database should persist and not be lost.

Special Thanks

- This course is originally taught by Dr. T. M. Murali
  - I am using Dr. Murali's course material