



Introduction to Database Systems



CS432

Instructor
Jayavel Shanmugasundaram



CS432/433: Introduction to Database Systems

- ❖ How does Wal-Mart manage its 200 TB data warehouse?
- ❖ What is the database technology behind ebay's website?
- ❖ How do you build an Oracle 9i, IBM DB2 or Microsoft SQL Server database?



CS432/433: Introduction to Database Systems

- ❖ Underlying theme: How do I build a database system?
- ❖ CS432 will deal with the underlying *concepts*
 - No programming assignments
 - Big change from previous years!
- ❖ CS433 will be the *practicum*
 - Build components of "realistic database system" (C++ programming)



CS432 Course Information

- ❖ Information is one of the most valuable resources in this information age
- ❖ How do we effectively and efficiently manage this information?
 - Relational database management systems
 - ◆ Dominant data management paradigm today
 - 6 billion dollar a year industry
 - ◆ You will see this in the job market!



Prerequisites

- ❖ Courses
 - CS212 (Computers and Programming)
 - CS312 (Structure and Interpretation of Computer Programs)



People

- ❖ Instructor
 - Jayavel Shanmugasundaram
- ❖ TAs
 - Adina Crainiceanu
 - Adrian Munteanu
 - Warren Wong
 - Lin Zhu

Access to Instructor and TAs

- ❖ Office hours
 - Posted on course web site
 - <http://www.cs.cornell.edu/courses/cs432>
- ❖ Course newsgroup
 - Monitored by TAs
 - Reply within 24 hours on weekdays, 48 hours on weekends
- ❖ TA mailing list
 - cs432ta@cs.cornell.edu
 - Do not directly email TAs

Course Structure

- ❖ Three components
 - Class lectures (5% for class participation)
 - Assignments (40%)
 - Examinations (55%)
- ❖ No programming assignments in CS432
 - Big change from previous years!
 - CS433 will have all programming assignments

Class Lectures

- ❖ Textbook: "Database Management Systems"
 - By Raghu Ramakrishnan and Johannes Gehrke
 - Required textbook
- ❖ Syllabus
 - Defined by class lectures
 - Not defined by textbook

Course Structure

- ❖ Three components
 - Class lectures (5% for class participation)
 - Assignments (40%)
 - Examinations (55%)

Assignments

- ❖ Eight assignments
 - Seven written assignments
 - One SQL assignment
- ❖ Each assignment worth 5% of total grade

Assignment Policies

- ❖ Assignments have to be done individually
 - No collaboration with others
- ❖ Academic integrity violations taken VERY seriously
 - Read Cornell and CS academic integrity policies
 - Available off course web page
 - Need to sign and hand in form
- ❖ Course management system used to post assignment grades

Assignment Policies (contd.)

- ❖ No late submissions
 - Will receive 0% of grade for late submissions
 - No exceptions (assignments handed out well in advance of deadline)
- ❖ Regrade requests
 - Within 7 days after assignments are graded
 - Hard deadline

Course Structure

- ❖ Three components
 - Class lectures (0%, but attendance is crucial ☺)
 - Assignments (60%)
 - Examinations (40%)

Exams

- ❖ Mid-term exam (20%)
 - 23 October 2003, 7:30-9:30pm
 - Closed book exam
- ❖ Final exam (35%)
 - Examination period
 - Closed book exam
 - Cumulative with emphasis on second half
- ❖ Do not schedule other events on these days

Relationship to CS433

- ❖ CS432 is about *concepts* underlying databases
 - No programming assignments
- ❖ CS433 is the *practicum* associated with CS432
 - Will actually build a “realistic” database system
 - C++ programming (okay if you know Java)
- ❖ Complementary
 - Suggest that you take both
 - **Can** take CS432 without taking CS433
 - **Cannot** take CS433 without taking CS432

Is CS432/433 a lot of work?

- ❖ It depends!
 - Much of the material in CS432 is probably new to you
 - CS433 has substantial programming assignments
- ❖ Then why on earth should I take this course?
 - Intellectual argument
 - ◆ Big conceptual ideas
 - ◆ Meeting of theory and practice
 - Utilitarian argument
 - ◆ Many, many real applications (digital libraries, web, ...)
 - ◆ Job market!
 - Others have gone through worse
 - ◆ CS432/433 were a single course before!

CS530: Architecture of Large-Scale Information Systems

- ❖ **How do you build e-commerce websites such as amazon.com?**
- ❖ **How do you build a reliable service that scales to millions of users?**
- ❖ **How are Internet transactions processed?**
- ❖ **How do you manage audio, video and XML data?**

CS530: Architecture of Large-Scale Information Systems

- ❖ Underlying theme: How do I build *applications* on top of a database system?
- ❖ Will combine coverage of fundamental concepts with “hands-on” experience
- ❖ Prerequisite: CS432

CS530: Material Covered

- ❖ Three-tier architectures
- ❖ Edge caches
- ❖ Distributed transaction management
- ❖ Web services
- ❖ Content management

- ❖ Technologies: .NET, JSPs, ASPs, Servlets, Enterprise Java Beans (EJBs), XML, SOAP

Reminder

- ❖ Complete academic integrity form
 - Need to hand this in for your course management system account