

KSU
IS Department
IS533 - Advanced Topics in Databases
Fall 2011/2012

Instructor: Dr. Bassam Hammo (Associate Prof. of Natural Language Processing)

Office: IS 2043

Email: bhammo@ksu.edu.sa

Course homepage: *staff.ksu.edu.sa/bhammo*

Class meetings: on Tuesday (10:30-1:20) Studio#5

Office Phone Number: 4698725

Contact me: via email with IS 533 in the subject line. I'll try my best to get back to you within 24 hours.

Course mailing list: An **email list** for this course will be compiled soon. You need to subscribe yourself to this list ASAP in order to receive messages related to this course.

ILOs (Intended Learning Outcomes):

This course will provide you with an overview of a selected set of the advanced topics in database systems. The goal is to expose you to active areas in databases by discussing book chapters in the classroom.

Course Outline:

Database Systems: Semantic data modeling, Object-oriented databases, Query optimization, Semantic Integrity. Distributed Databases (DDB), Data fragmentation and distributed transparency, Distributed query processing, Concurrency control methods: Serializability in a DDB and the two-phase locking method, Concurrency control based on timestamps, Recovery management.

You will also conduct a research topic in an area of your choice from the recent literature related to the database area. You are required to present your topic in front of the class (A suggested list of topics will be provided to you).

Text Book:

There will *not* be an assigned textbook in the course. However, we will discuss *selected chapters* from several books and papers from the literature. I will link the material onto our course **webpage** frequently.

Optional Readings:

- Fundamentals of Database Systems, 5th Edition or later, R. Elmasri, and S. Navathe, Benjamin Cummings.
- Database System Concepts, 5th Edition or later, H.F. Korth and A. Silberschatz, McGraw-Hill.

- Principles of Distributed Database Systems, 2nd Edition **or later**, M. Tamer Ozsu, Patrick Valduriez.

Topics to be covered (based on time):

There will be a core collection of readings in selected topics in databases around which our discussions will be centered. Topics cover may include a subset of the following:

- Database Normalization and Functional Dependencies
- Database Storage & Recovery
- Transaction Management and Concurrency Control Protocols
 - 2PL Protocol
 - Timestamp Protocol
- Distributed Databases
 - Parallel Management Distribution
 - Data Distribution
 - Fragmentation: Horizontal & Vertical
 - Allocation
 - Replication
 - Heterogeneous Databases and Data Integration.
- Query Optimization
- XML Databases
- Temporal and Spatial Databases
- Object-oriented/object-relational databases.
- Data warehousing.
- Database Mining and Knowledge Discovery.

Paper Presentations and Discussion:

A group of students (2-3) will be formed and assigned a topic to read and present during the semester. When you present a topic, you should meet the following goal:

- Provide a background of the topic you selected from the literature (ACM > 2008). Why is this topic important? What problem(s) it solves? In some cases, you may find it useful to explain how the concepts in the selected topic advance the state of the art.

Grading Policy:

Final grades are computed based on 100 points:

- 10 points: Active Participation & in Class Discussions.
- 20 points: Topic Presentation & Discussion
- 30 points: Midterm Exam
- 40 points: Final Exam