

The University of Jordan/ KASIT
CIS Department
DATABASE SYSTEMS
1902723
Second Semester 2014/2015
Advance Topics in Database Management Systems

Instructor

- Omar Adwan, Assistant Professor
- Email: adwanoy@ju.edu.jo
- Office Phone Number: 22610

Course Information

- Course homepage: elearning.ju.edu.jo
- Class time : Mon.,Wed. [2:00-3:30]
- Prerequisites: Undergraduate database course.

Course Description

This course is to introduce graduate students with an advance topics in database systems, which are relevant to research and industrial development going on today in the broad context of database system. The focus of this course is on normalization and FD's, transaction management and concurrency control: serializability in DB and the two-phase locking method, timestamps and recovery techniques. This course brings together several key information technologies used in manipulating, storing, and analyzing big data. We review MapReduce techniques for parallel processing and Hadoop, an open source framework. We touch on related tools that provide SQL-like access to unstructured data. We analyze so-called NoSQL storage solutions. We work with the public cloud as unlimited resource for big data analytics.

Classes will consist of lectures and discussions based on readings from the database literature.

Course Textbooks

We will be discussing selected chapters from several different books and papers from the literature. I will link the material onto our course webpage.

The course readings will primarily be drawn from the following books:

1. Database Systems Concepts, 6th Edition by Silberschatz, Korth, and Sudarshan. ISBN: 0073523321.
2. Fundamentals of Database Systems, 6th Edition by Ramez Elmasri, Shamkant B. Navathe. ISBN-13: 978-0-136-08620-8
3. Hellerstein, Joseph M., and Michael Stonebraker, eds. Readings in Database Systems. 4th ed. Cambridge, MA:MIT Press, 2005. ISBN: 0262693143.
4. M. Tamer Özsu and Patrick Valduriez, Principles of Distributed Database Systems. 2nd ed, Prentice Hall, 1999.
5. Ramakrishnan, Raghu, and Johannes Gehrke. Database Management Systems. 2nd ed. New York, NY: McGraw-Hill, 2000. ISBN: 0072465352.
6. Ramez Elmasri, Shamkant B. Navathe. Fundamentals of Database Systems. 5th Edition, Addison-Wesley Publishing

Learning Outcomes

Upon completing this course, students should:

1. Be able to discuss/explain the importance of new data modeling concepts.
2. Be able to apply efficient and effective design techniques of new data modeling concepts.
3. Understand and apply transaction management mechanism and its related issues.
4. Understand recovery control protocols.
5. Understand and apply advanced SQL constructs.
6. Understand database security and integrity techniques for modern databases.
7. Understand technologies used in manipulating, storing, and analyzing big data and related issues.

Core Topics

The course will cover the following topics:

No.	Topics
1	Database Design <ul style="list-style-type: none">• Database Normalization• Functional Dependencies Theory• Decomposition• More Normal Forms
2	Data Storage and Querying <ul style="list-style-type: none">• Physical Storage Media• Records Organization
3	Indexing and Hashing <ul style="list-style-type: none">• Indexing Concepts• Ordered Indices• B+- Tree• Hashing• Bitmap Indices
4	Transaction Management <ul style="list-style-type: none">• Transaction and System Concepts• Desirable Properties of Transactions• Characterizing Schedules Based on Recoverability• Characterizing Schedules Based on Serialization
5	Concurrency Control Protocols <ul style="list-style-type: none">• Lock-Based Protocols• Deadlock Handling• Multiple Granularity• Timestamp-Based Protocols

5	Recovery System <ul style="list-style-type: none"> • Recovery Concepts • Recovery Techniques Based on Deferred Update • Recovery Techniques Based on Immediate Updates
6	Big Data Analytics <ul style="list-style-type: none"> • Big Data Concepts • Current challenges, trends, and applications • Hadoop • unstructured data concepts • Map Reduce technology • NOSQL • Public Cloud as a resource for big data analytics
7	Distributed Databases <ul style="list-style-type: none"> • Parallel Management Distribution • Data Distribution • Fragmentation: Horizontal & Vertical • Allocation • Replication • Heterogeneous Databases and Data Integration

Course objectives:

This course will provide students with an overview of a selected set of the advanced topics in database systems. At the end of this course, the student will become familiar with the fundamental concepts of Big Data management and analytics; will become competent in recognizing challenges faced by applications dealing with very large volumes of data.

Final project

Projects are one of the most important learning tools of this class. The final project is entirely to the discretion of the student (upon instructor approval). Students are free to explore a problem of their interest and propose their own solution. The project has the following deliverables:

- **Proposal.** Maximum 1 page of project proposal, why the problem is important, what has been done so far in the field, and what are the expected outcomes
- **Final report.** Maximum 10 page report highlighting consisting on the traditional sections of a term paper includes: Abstract, introduction, Literature Review, motivation, method, results, and conclusion

Projects will be done individually.

POLICIES

Class attendance:

Attendance to class is expected and note taking encouraged. Important information (about assignments, projects, policies) may be communicated only in the lectures. We may also cover additional material (not available in the notes) during the lecture.

Participation

Participation is the barometer of the class. Based on it I can determine if the pace of the course is too fast or too slow, it helps me to spot pitfalls and misconceptions, and it helps you to reinforce the material you learned.

Participation accounts for your final grade and won't be given for granted. You are required to participate either in class or electronically.

Feedback:

I value student's opinions regarding the course and I will take them in consideration to make this course as exciting and engaging as possible. Thus, through the semester I will ask students formal and informal feedback. Formal feedback includes short surveys on my teaching effectiveness, preferred teaching methods, and pace of the class. Informal feedback will be in the form of polls or in-class questions regarding learning preferences. Remember that it is in the best interest of the class if you bring up to my attention if something is not working properly (e.g the pace of the class is too slow, the projects are boring, my teaching style is not effective) so that I can make the corrective steps.

Grading

Activity	Percentage
Topic Presentation & Discussion	10%
Writing a research paper	20%
Midterm Examination	30%
Final Examination.	40%