COSC6340: Database Systems

Instructor: Carlos Ordonez

1 Course information

Schedule: MoWe 1pm-2:30pm
Google newsgroup: "COSC6340-". Instructions on TA web page.
email: lastname AT cs uh edu (Start subject line with "COSC6340-")

2 Course contents

This is a graduate level course on database management systems (DBMSs). The textbook is [2], complemented by [3] and [1]. The course will require reading some research papers electronically available.

Topics include the following. Fundamental theory: first order logic, relational model and relational algebra. Dependencies, normalization up to 5NF. SQL: SPJA queries, derived tables/view, pivoting, cube/pivot, recursive queries, UDFs, stored procedures. Database design and software engineering: ER model, workflows. Internal subsystems of DBMS: secondary storage, main storage architectures, buffer management, indexing data structures, concurrency control, transaction processing, query optimizer, fault tolerance: for transaction processing (recovery) and for long query processing (MPP). Overview of security, cube and big data analytics.

3 Grading

- 70%: 2 programming projects.
- 30%: Midterm (2/3 of course) exam.

Project 1 will involve normalization and transaction processing. Project 2 will involve creating an SQL query generator to perform analytics on a large database via SQL queries; Project 1 will be delivered will take 5 weeks and Project 2 will take 7 weeks. DBMSs will work on Unix (Linux Ubuntu). Programming will be done in SQL, Java and C++. Programming assignments must be done in pairs (i.e. a team of 2 students), that will be assigned by instructor. Programs will be carefully tested by TA for correctness and efficiency. The best projects will be shown to the Instructor.

References

[1] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases : The Logical Level*. Pearson Education POD, facsimile edition, 1994.

[2] H. Garcia-Molina, J.D. Ullman, and J. Widom. *Database Systems: The Complete Book*. Prentice Hall, 1st edition, 2001.

[3] J. Han and M. Kamber. *Data Mining: Concepts and Techniques*. Morgan Kaufmann, San Francisco, 1st edition, 2001.