SYLLABUS

IDS 521: Advanced Database Management
Call# 16760 (4 credit hours) – Spring 2017 (Prerequisite: IDS 410)

Instructor: Yann Chang, Ph.D.
Office/Phone/email: Room 2126 UH; (312) 355-1245; jsun@uic.edu
Class Time: Saturday 12:15 P.M. – 3:15 P.M.
Class Location: Douglas Hall room 210
Office Hours: Tuesday and Thursday from 10:30AM to 12:00PM or by appointment

Course Introduction and Overview

In general, we can categorize Information Technology systems into an online transactional processing (OLTP) system and an online analytical processing (OLAP) system. An OLTP system is characterized by a large number of short online transactions (READ, INSERT, UPDATE, and DELETE). The purpose of this system is to run fundamental business tasks. It processes transactions in a multi-user environment. It captures and stores detailed records of transactions in an OLTP database, which is typically based on the Relational data model and is normalized to the third normal form (3NF) with many tables. In contrast, an OLAP system is characterized by relatively low volume of transactions. Queries are often very complex and involve aggregations. The purpose of an OLAP system is to help decision makers with planning, problem solving, performance analysis, and data mining. An OLAP database contains accumulation of historical transactional data, aggregated measures, and multiple dimensions. An OLAP database is typically based on the dimensional data model and is de-normalized (usually star schemas or OLAP cubes) with fewer tables. In most OLTP and OLAP systems, the underlying database systems are based on the Relational data model, which has been around since early 1970. Starting around 2008, an explosion of new database systems occurred; and none of these systems adhered to the traditional relational implementations. These new databases, also known as NoSQL databases, were designed to store and process an exponentially growing quantity of data ("Big Data"). These NoSQL databases pose both opportunities and challenges to OLTP and OLAP systems.

In this course, we will discuss three different but related topics: (a) OLTP database design, SQL, indexing, and query evaluation; (b) OLAP dimensional database design, Extract-Transform-Load process, OLAP cubes; and (c) NoSQL databases. For the first five sessions, we will quickly review some materials (Entity Relationship model and Relational data model) from the prerequisite course (IDS 410). Then, we will discuss Relational Algebra, SQL, indexing, and query evaluation. For the next five sessions, we will focus on OLAP database (Data Warehouse or Data Mart) design using dimensional modeling. We will discuss fact tables for measurements, dimension tables for providing descriptive contexts to the fact tables, the Extract-Transform-Load process, and the OLAP cubes. We will use Microsoft SQL Server and Analysis Services running on SQL Server platform to illustrate the dimensional and multidimensional modeling concepts. Finally, for the last five sessions, we will discuss different categories of NoSQL databases.

Course Objectives

To review the basic concepts of Entity Relationship model, Relational data model, data normalization, and SQL. To introduce students to the Relational Algebra, indexing, query evaluation, dimensional modeling, and NoSQL databases. To provide students with conceptual skills to design both OLTP and OLAP databases.

Assignments, Quizzes, and Exams

For the first topic (OLTP database design), there will be two assignments (5% each), one online quiz (5%), and one exam (15%). For the second topic (OLAP database design), there will be two assignments (5% each), two online quizzes (5% each), and one exam (20%). For the third topic (NoSQL databases), there will be one online quiz (5%), one project presentation (5%), and a group project (20%).

Letter Grades Determination

A – Total score of 90% or more
B – Total score of 80% and below 90%
C – Total score of 70% and below 80%
D – Total score of 60% and below 70%
F – Total score of less than 60%

Total Score Calculation

Total score = 0.05*(Hw1+Hwk2) + 0.05*Quiz1 + 0.15*(Exam1) + 0.05*(Hw3+Hw4) + 0.05*(Quiz2+Quiz3) + 0.20*(Exam2) + 0.05*Quiz4 + 0.05*Presentation + 0.20*(Project)

Depending on class performance, the grades might be curved.

Three Required Texts:

Title: Database Management Systems, 3rd Edition
Authors: Raghu Ramakrishnan and Johannes Gehrke
Publisher: McGraw-Hill
Year: August 2002
ISBN: 978-0072465631

Title: The Data Warehouse Toolkit: The Definitive Guide to Dimensional Modeling, 3rd edition
Authors: Ralph Kimball and Margy Ross
Publisher: John Wiley & Sons
Year: July 2013
ISBN: 978-1118530801

Harrison, G., & SpringerLink (Online service). (2015). Next generation databases NoSQL, NewSQL, and Big Data. Berkeley, CA: Apress. doi:10.1007/978-1-4842-1329-2
This book is available free online from the UIC library: http://bit.ly/2bF84m4

Recommended (optional) Text:

Title: Delivering Business Intelligence Solutions with Microsoft SQL Server 2012, 3rd edition
Author: Brian Larson
Publisher: McGraw Hill
Year: 2012
ISBN: 978-0-07-175938-0

Tentative outline

Date	Subject	Reading Assignments
Week 1	Introduction to Database Design; The Relational Model	Ramakrishnan and Gerhke: Chapters 2 and 3
Week 2	Relational Algebra; SQL	Ramakrishnan and Gerhke: Chapters 4 and 5
Week 3	Overview of Storage and Indexing	Ramakrishnan and Gerhke: Chapter 8
Week 4	Overview of Query Evaluation	Ramakrishnan and Gerhke: Chapter 12
Week 5	**Review for Exam 1 (60 minutes) and then Exam 1 (90 minutes)**	
Week 6	Data Warehousing, Business Intelligence, and Dimensional Modeling Primer; Kimball Dimensional Modeling Techniques Overview	Kimball and Ross: Chapters 1 and 2
Week 7	Retail Sales; Inventory	Kimball and Ross: Chapters 3 and 4
Week 8	Procurement; Order Management	Kimball and Ross: Chapters 5 and 6
Week 9	Accounting	Kimball and Ross: Chapter 7
Week 10	**Review for Exam 2 (60 minutes) and then Exam 2 (90 minutes)**	Kimball and Ross: Chapters 1-7
Week 11	Next Generation Databases	Harrison: Chapters 1-2
Week 12	Next Generation Databases (continued)	Harrison: Chapters 3-4
Week 13	Next Generation Databases (continued)	Harrison: Chapters 6-7
Week 14	**Group Project Presentation – Part 1**	
Week 15	**Group Project Presentation – Part 2**	
Week 16	**No final exam**	

College of Business Administration
University of Illinois at Chicago

January 2003

Honor Code for the College of Business Administration

As an academic community the College of Business Administration at the University of Illinois at Chicago is committed to providing an environment in which teaching, learning, research, and scholarship can flourish and in which all endeavors are guided by academic and professional integrity. All members of the college community – students, faculty, staff, and administrators – share the responsibility of insuring that high standards of integrity are upheld so that such an environment exists.

In pursuit of these high ideas and standards of academic life, as a student I hereby commit myself to respect and uphold the University of Illinois at Chicago (UIC) College of Business Administration Honor Code during my entire matriculation at UIC. I agree to maintain the highest moral and ethical standards in all academic and business endeavors and to conduct myself honorably as a responsible member of the college academic community. This includes the following:

- Not to seek unfair advantage over other students, including, but not limited to giving or receiving unauthorized aid during completion of academic requirements;
- To represent fact and self truthfully at all times;
- To respect the property and personal rights of all members of the academic community.

An abbreviated version of the Honor Code pledge may be printed on course syllabi, exam booklets, and other uses as deemed appropriate. The abbreviated version is as follows:

> This course and its associated coursework are being administered under the policies of the University of Illinois at Chicago (UIC) College of Business Administration Honor Code. All students are expected to respect and uphold this code.

Honor Code Violations and Enforcement

Violations of the Honor Code are just causes for discipline under the University of Illinois at Chicago Student Disciplinary Procedure, and all allegations of Honor Code violations shall be handled pursuant to that Procedure.
(For a complete description of just causes for discipline, disciplinary procedures, and sanctions, see the pamphlet “Student Disciplinary Procedure of the Senate Committee on Student Discipline,” available from the Office of the Dean of Students, SSB, Suite 3030, 1200 W. Harrison St., M/C 318.)

Honor Code Council

The Honor Code Council shall be formed consisting of two faculty members elected by faculty vote and six student members (two graduate students and four undergraduate students) appointed by the Assistant Dean for Student Services through recommendation of the Dean’s Advisory Council (DAC). Each member will serve a term of one academic year, beginning in August. The Council shall elect a chairperson and a secretary.

The duties of the Honor Code Council shall include:

- Recommend changes in the Honor Code.
- Inform all students, staff, and faculty of the Honor Code and the procedures to be followed for pursuing alleged Honor Code violations.
- Ensure that the Honor Code is prominently displayed within the College of Business Administration and on course syllabi and exam booklets.
- Ensure that the Honor Code and related information are accurately described in the Graduate and Undergraduate catalogs.
- Work with the Office of Student Services to include the Honor Code in its promotion of the University of Illinois at Chicago College of Business Administration to potential students.
- Work with college administration to ensure that the Honor Code statement is signed by all students prior to their enrollment in the college.
- Inform all UIC faculty, staff, and students of the Honor Code of the College of Business Administration, and encourage the adoption of its principles.

A database management system (DBMS), sometimes just called a database manager, is a program that lets one or more computer users create and access data in a database.Â Functions of a DBMS. The DBMS manages three important things: the data, the database engine that allows data to be accessed, locked and modified, and the database schema, which defines the database's logical structure. These three foundational elements help provide concurrency, security, data integrity and uniform data administration procedures.