

COSC 404 - Database System Implementation

Winter 2007 Term 2

Instructor: Dr. Ramon Lawrence
Class Schedule: 9:30 a.m. – 11:00 a.m. Tuesday/Thursday
Location: LIB 303
Lab time/location: 10:30 a.m. – 12:30 p.m. Wednesdays at SCI 234
Office Hours: 1:00–3:00 p.m. Monday and 1:00-3:00 p.m. Thursday or by appointment
Office Location: SCI 263
Phone: 807-9390
E-mail: ramon.lawrence@ubc.ca (preferred contact method)
Course URL: <http://people.ok.ubc.ca/rlawrenc/404/>

Course Description

Official Calendar: Review of database environment and database design principles. Recovery and concurrency control in distributed database systems, object and object-relational databases, data mining, and data warehousing. Client/server-based database application is designed and developed using state-of-the-art technology. OUC equivalent: COSC 404.

Specific description: This course covers advanced database implementation and design topics including file organizations, storage management, database system architectures, query optimization, transaction management, recovery, and concurrency control. Additional topics including distributed databases, mobile databases, and integration may also be covered.

Prerequisite

- COSC 304 – Introduction to Database Systems

Marking and Evaluation

Assignments	20 % (approximately 5-6 assignments)
Midterm Exam	20 % (75 minutes in class)
Final Exam	40 % (cumulative, two hours)
Project	20 %

Textbook and Reference Material (*Optional*)

Optional textbook: Hector Garcia-Molina, Jeffrey Ullman and Jennifer Widom, *Database Systems: The Complete Book*, Prentice Hall, ISBN 0-130-31995-3, 2002.

Expectations

- I expect students to attend **all** classes and prepare before attending class. This includes reading relevant sections of the textbook and reviewing notes from previous lectures.
- I recommend all students download and read a copy of the lecture notes **before** the lecture.
- I expect all students to learn the material in the course and undertake sufficient effort to produce all the programming assignments and quality projects.
- I want all students to enjoy attending class and feel free to participate according to their own personalities. Feel free to ask questions by raising your hand or speaking out at appropriate times.
- Please actively participate in class discussions, questions, and problem solving exercises.
- **I want all students to pass the course, receive a good grade, and feel the course was beneficial.**

Homework Expectation

For this course, it is expected that you will spend *at least six hours per week in out-of-class preparation*.

Grievances and Complaints Procedures

A student who has a complaint related to this course should follow the procedures summarized below.

- The student should attempt to resolve the matter with the instructor first. Students may talk first to someone other than the instructor if they do not feel, for whatever reason, that they can directly approach the instructor.
- If the complaint is not resolved to the student's satisfaction, the student should go to the departmental chair Cynthia Mathieson, at ART 300, 807-8730.

Your Responsibilities

Your responsibilities to this class and to your education as a whole include attendance and participation. You have a responsibility to help create a classroom environment where all may learn. At the most basic level, this means you will respect the other members of the class and the instructor and treat them with the courtesy you hope to receive in return. Inappropriate classroom behavior may include: disruption of the classroom atmosphere, engaging in non-class activities, talking on a cell-phone, inappropriate use of profanity in classroom discussion, use of abusive or disrespectful language toward the instructor, a student in the class, or about other individuals or groups.

Academic Dishonesty

The academic enterprise is founded on honesty, civility, and integrity. As members of this enterprise, all students are expected to know, understand, and follow the codes of conduct regarding academic integrity. At the most basic level, this means submitting only original work done by you and acknowledging all sources of information or ideas and attributing them to others as required. This also means you should not cheat, copy, or mislead others about what is your work.

Violations of academic integrity (i.e., misconduct) lead to the break down of the academic enterprise, and therefore serious consequences arise and harsh sanctions are imposed. For example, incidences of plagiarism or cheating usually result in a failing grade or mark of zero on the assignment or in the course. Careful records are kept in order to monitor and prevent recidivism. A more detailed description of academic integrity, including the policies and procedures, may be found at:

<http://web.ubc.ca/okanagan/faculties/resources/academicintegrity.html>

If you have any questions about how academic integrity applies to this course, please consult with your professor.

Disability Assistance

If you require disability-related accommodations to meet the course objectives, please contact the Coordinator of Disability Resources located in the Student Development and Advising area of the student services building. For more information about Disability Resources or academic accommodations, please visit: <http://okanagan.students.ubc.ca/current/disres.cfm>

Missing an Exam

Only students who miss an exam for a reason that corresponds to the University of British Columbia Okanagan's policy on excused absences from examinations will be permitted to take the exam at a later time. A make-up exam may have a question format different from the regular exam. If the reason for absence is satisfactory, the student may either take the exam, or if a midterm exam is missed, the student's final exam will be worth more of the final grade.

Course Outline

The course has a substantial amount of material to be covered in a short time. This requires the student make a strong effort to keep up with the material discussed in class. Below is an outline of the topics. The professor is not bound to these topics and timelines as they only serve as a general reference.

Date	Topics Covered and Description
January 8 (T)	First day of classes. Introduction to course, discuss syllabus/project, WebCT
January 10 (Th)	Storage issues I: memory hierarchy, hard drive technology, RAID
January 15 (T)	Storage issues II: file organizations, data representation, record types (fixed vs. variable)
January 17 (Th)	Storage issues III: storing records in blocks, file operations, buffering, pointer swizzling
January 22 (T)	Indexing I: motivation, index types, index maintenance, primary/secondary indexes
January 24 (Th)	Indexing II: B-Trees - 2-3 trees, B and B+-Trees
January 29 (T)	Indexing III: Hashing - main memory, external, extendible, linear
January 31 (Th)	Indexing IV: SQL indexing, multi-value indexing, spatial indexing using R-trees
February 5 (T)	Query processing I: components, query plans, basic scans, iterators, one-pass algorithms
February 7 (Th)	Query processing II: nested-loop joins, external sorting, two-pass algorithms (sorting)
February 12 (T)	Query processing III: sort-join, hash-join, two-pass algorithms (hashing), index join
February 14 (Th)	Query processing IV: programming iterators in Java
February 19 (T)	No classes during Midterm Break.
February 21 (Th)	No classes during Midterm Break.
February 26 (T)	Query optimization I: query parsing, relational algebra laws, heuristic query optimization
February 28 (Th)	Midterm Exam In Class.
March 4 (T)	Query optimization II: physical query plans, pipelining, cost-based query optimization
March 6 (Th)	Transaction processing I: overview, transaction states, ACID properties, schedules
March 11 (T)	Transaction processing II: conflict/view serializability, precedence graphs
March 13 (Th)	Concurrency control I: overview, locks, two-phase locking (2PL), graph protocols
March 18 (T)	No class.
March 20 (Th)	No class.
March 25 (T)	Concurrency control II: multi-granularity locking, timestamps and validation protocols
March 27 (Th)	Concurrency control III: deadlock handling, starvation, wait-for graphs
April 1 (T)	Recovery I: motivation, shadow paging, log-based recovery
April 3 (Th)	Recovery II: checkpoints, undo/redo logging, deferred versus immediate update
April 8 (T)	Project presentations.
April 10 (Th)	Final programming project due. Review for final exam.

Laboratory times: The laboratory time will be spent on written and programming lab assignments.

Week	Dates	Topics Covered and Description
1	January 9	No Lab First Week of Class
2	January 16	Lab 1: Storage Issues
3	January 23	Lab 1: Storage Issues (cont.)
4	January 30	Lab 2: Indexing
5	February 6	Lab 2: Indexing (cont.)
6	February 13	Lab 3: Query Processing
7	February 20	No Lab During Midterm Break
8	February 27	Lab 4: Query Optimization
9	March 5	Lab 4: Query Optimization (cont.)
10	March 12	Lab 5: Transactions
11	March 19	Lab 5: Transactions (cont.)
12	March 26	Lab 6: Concurrency Control and Recovery
13	April 2	Lab 6: Concurrency Control and Recovery (cont.)
14	April 9	No Lab Last Week of Class