

# COSC 404 / COSC 504 - Database System Implementation

## 2017 Winter Term 2

**Instructor:** Dr. Ramon Lawrence  
**Class Schedule:** 11:00 a.m. – 12:30 p.m. Tuesday/Thursday  
**Location:** EME 1202  
**Lab times/locations:** **L01:** 9:00 a.m. – 11:00 a.m. Thursdays at SCI 234  
**L02:** 12:00 p.m. – 2:00 p.m. Wednesdays at ART 215  
**Office Hours:** 9:00-10:00 a.m. Tuesdays/Thursdays, in labs, or by appointment  
**Office Location:** ASC 349  
**Phone:** 807-9390  
**E-mail:** [ramon.lawrence@ubc.ca](mailto:ramon.lawrence@ubc.ca) (preferred contact method)  
**Course URL:** <http://people.ok.ubc.ca/rlawrenc/teaching/404/>

### Course Description

**Official Calendar:** Fundamental concepts in constructing database systems including file organizations, storage management, system architectures, query processing/optimization, transaction management, recovery, and concurrency control. Additional topics may include distributed databases, mobile databases, and integration. [3-2-0]

### Prerequisite

- COSC 304 – Introduction to Database Systems (60% or higher)

### Evaluation Criteria and Grading

Programming Assignments	20 %	(weekly assignments)
Written Assignments	10 %	(weekly assignments)
Clickers	5 %	
Two Midterm Exams	30 %	
Final Exam	35 %	(cumulative, three hours)

### Graduate Student Evaluation:

Programming Assignments	10 %	(weekly assignments)
Written Assignments	5 %	(weekly assignments)
Clickers	5 %	
Project	15 %	
Two Midterm Exams	30 %	
Final Exam	35 %	(cumulative, three hours)

- Graduate students are responsible for a substantial research and development project. Optional or bonus assignment questions for undergraduates will often be required for graduate students.
- A student must receive a combined grade of at least 50% on the exams (midterms and final) to pass the course. Otherwise, the student has a maximum overall grade of 45.

### Textbook and Reference Material:

- A *clicker* is required. Notes are online and as a course pack available at the bookstore.
- A textbook is *not required*. Students can get supplemental material from any database textbook.

## Expectations

- Attend **all** classes and prepare before attending class.
- Read the lecture notes **before** the lecture.
- Learn the material in the course by completing all assignments.
- Enjoy attending class and feel free to participate according to your personality. 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, expect to 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 John Braun at SCI 388, 807-8032.

## 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 Integrity

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 breakdown 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 to monitor and prevent recidivism. A more detailed description of academic integrity, including the policies and procedures, may be found at <http://www.calendar.ubc.ca/okanagan/index.cfm?tree=3,54,111,959>. If you have any questions about how academic integrity applies to this course, consult with the instructor.

## Disability Assistance

If you require disability-related accommodations to meet the course objectives, please contact the Diversity Advisor of Disability Resources located in the University Centre, Room 227. For more information about Disability Resources or academic accommodations, please visit the website at: <http://students.ok.ubc.ca/drc/welcome.html>

## Equity, Human Rights, Discrimination and Harassment

UBC Okanagan is a place where every student, staff and faculty member should be able to study and work in an environment that is free from human rights based discrimination and harassment. If you require assistance related to an issue of equity, discrimination or harassment, please contact the Equity Office, your administrative head of unit, and/or your unit's equity representative. **UBC Okanagan Equity Advisor: ph. 250-807-9291; email [equity.ubco@ubc.ca](mailto:equity.ubco@ubc.ca)**  
**Web:** [www.ubc.ca/okanagan/equity](http://www.ubc.ca/okanagan/equity)

## Missing an Exam

Only students who miss the final 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 final exam at a later time. A make-up exam may have a question format different from the regular exam. **There will be no make-up midterm exams.** If the reason for absence is satisfactory, the student's final exam will be worth more of the final grade. Further information on Academic Concession can be found under Policies and Regulation in the Okanagan Academic Calendar <http://www.calendar.ubc.ca/okanagan/index.cfm?tree=3,48,0,0>.

## Objectives and Learning Outcomes

- Experience using and developing on many different SQL and NoSQL databases.
- Manipulate data in memory and storage and use index structures for improved performance.
- Understand query processing including parsing, translation, optimization, and execution.
- Apply principles of transactions, concurrency, recovery, and distribution for databases.
- Use knowledge of database techniques to be better users with the ability to use different database systems, compare their properties, and adapt database techniques when developing software.

## 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 prepare before class so that the material can be practiced 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 4 (TH)	<b>First day of classes. Introduction to course. Database architecture.</b>
January 9 (T)	Storage: Accessing/Representing Data on Devices (SSD,RAID) using Records/Blocks
January 11 (TH)	Indexing I: Index Types, Primary Indexes, Multi-level Indexes, Secondary Indexes
January 16 (T)	Indexing II: B-Trees (insertion, deletion), B+-Trees
January 18 (TH)	Indexing III: B+-Trees, R-Trees
January 23 (T)	Indexing IV: Hash Indexes, SQL Indexing in Practice
January 25 (TH)	Query processing I: SQL/RA Review, Types of Operators, Iterators, One-pass Algorithms
January 30 (T)	Query processing II: Nested-Loop Joins, External Sorting, Two-Pass Sorting Algorithms, Sort-Join, Sort-Merge-Join
February 1 (TH)	Query processing III: Hash Partitioning, Two-Pass Hash Algorithms, Hybrid Hash Join
February 6 (T)	Query optimization I: Query Parsing/Translation, Relational Algebra Laws
February 8 (TH)	<b>Midterm exam #1</b>
February 13 (T)	<b>No classes during Midterm Break.</b>
February 15 (TH)	<b>No classes during Midterm Break.</b>
February 20 (T)	Query optimization II: Heuristic Optimization, Physical Query Plans
February 22 (TH)	Query optimization III: Cost-based Query Optimization
February 27 (T)	Transaction processing I: ACID Properties, Schedules, Conflict Serializability
March 1 (TH)	Transaction processing II: View Serializability, Schedule Properties
March 6 (T)	Concurrency control I: Two-Phase Locking (2PL), Multiple Granularity Locking, Deadlock Handling, Wait-for Graphs
March 8 (TH)	Concurrency control II: Timestamp Protocols, Validation Protocols, Multi-versioning, Snapshot isolation
March 13 (T)	<b>Midterm exam #2</b>
March 15 (TH)	Concurrency control III: SQL Isolation Levels, Phantom Phenomenon, CC in systems
March 20 (T)	Recovery I: Types of Failures, Log-Based Recovery
March 22 (TH)	Recovery II: Undo/Redo Logging
March 27 (T)	Distribution I: Architectures, Semi-joins, Two-Phase Commit
March 29 (TH)	Distribution II: Fragmentation, Partitioning, Sharding
April 3 (T)	Distribution III: Replication: Master-Master and Master-Slave, CAP Theorem
April 5 (TH)	Architecture I: Comparison of database architectures: Relational, Key-Value, In-Memory <b>Review for final exam.</b>

**Laboratory times:** The laboratory time will be spent on programming assignments.

Week	Dates	Topics Covered and Description
1	January 1 - 5	<b>No Lab First Week of Class</b>
2	January 8 – 12	Lab 1: MySQL vs. PostgreSQL – Creating and Querying Data
3	January 15 – 19	Lab 2: MySQL vs. PostgreSQL – Indexing for Performance
4	January 22 – 26	Lab 3: Implementing a Text Database and JDBC Driver
5	Jan. 29 – Feb. 2	Lab 4: Query Processing with Iterators
6	February 5 – 9	Lab 5: Query Parsing with JavaCC
7	February 12 - 16	<b>No Lab During Midterm Break</b>
8	February 19 – 23	Lab 6: Storing JSON Documents: MongoDB and PostgreSQL
9	Feb. 26 – Mar. 2	Lab 7: Map-Reduce using MongoDB
10	March 5 – 9	Lab 8: Transactions with Microsoft SQL Server
11	March 12 – 16	Lab 9: VoltDB – In-memory database
12	March 19 – 23	Lab 10: Recovering from a Database Failure
13	March 26 – 30	Lab 11: Scaling MySQL: Master-Slave Replication and Partitioning
14	April 2 - 6	<b>No Lab Last Week of Class</b>

**Written assignments:** Written assignments practice fundamental skills. They will be done at various times including in class, in lab, and before and after class time.

Week	Dates	Topics Covered and Description
1	January 1 - 5	<b>No Written Assignment First Week of Class</b>
2	January 8 – 12	Assign 1: Storage and Index Performance Calculations
3	January 15 – 19	Assign 2: B-Trees/B+-Trees and Linear Hashing
4	January 22 – 26	Assign 3: Query Processing
5	Jan. 29 – Feb. 2	Assign 3: Query Processing (cont.)
6	February 5 – 9	Assign 4: Query Optimization
7	February 12 - 16	<b>No Assignment During Midterm Break</b>
8	February 19 – 23	Assign 4: Query Optimization (cont.)
9	Feb. 26 – Mar. 2	Assign 5: Transactions
10	March 5 – 9	Assign 6: Concurrency Control Protocols (2PL) and Deadlock Handling
11	March 12 – 16	Assign 7: Concurrency Control Protocols (timestamps)
12	March 19 – 23	Assign 8: Recovery
13	March 26 – 30	Assign 9: Distribution and Partitioning
14	April 2 - 6	<b>No Assignment Last Week of Class</b>