22C:244 – Database System Implementation
Spring 2004

Instructor:	**Dr. Ramon Lawrence**
Class Schedule:	1:05 – 2:20 p.m. Tuesday/Thursday
Location:	103 North Hall (and by video streaming)
Office Hours:	11:00–12:00 a.m. Tues/Thurs and 3:00-4:00 p.m. Tuesday or by appointment
Office Location:	201L MacLean Hall
Phone:	335-0561
E-mail:	ramon-lawrence@uiowa.edu (preferred contact method)
Course URL:	http://www.cs.uiowa.edu/~rlawrenc/teaching/244/index.html

Course Description

This course expands on 22C:144 to cover 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. A major component of the course is a database implementation project using current database languages and systems.

Prerequisite

- 21:124/22C:144 Introduction to Database Systems is recommended but not required

Marking and Evaluation

Homework Assignments		**10 %** (approximately 4 assignments)
Midterm Exam		**20 %** (75 minutes in class)
Final Exam		**40 %** (cumulative, two hours)
Programming Project		**30 %**
Proposal	5 %	
Presentation	5 %	
Final Program	20 %	

This course will use +/- grading. No late projects will be accepted.

Textbooks and Reference Material

- Recommended textbook: Hector Garcia-Molina, Jeffrey Ullman and Jennifer Widom, *Database Systems: The Complete Book*, Prentice Hall, ISBN 0-130-31995-3, 2002
- Alternate textbook: Thomas Connolly and Carolyn Begg, *Database Systems: A Practical Approach to Design, Implementation, and Management,* Addison Wesley, 3rd ed., ISBN 0-201-70857-4, 2002
- The textbooks are **optional**, although it is recommended you acquire one for reference. Most database textbooks have the required material, so you may use alternate textbooks as well.

Teaching Assistant

Eduard Dragut, eduard-dragut@uiowa.edu, Office: 201N MacLean Hall
Office Hours: **TBA**

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 advanced projects associated with a class of this level.
- I want all students to enjoy attending class and feel free to participate according to their own personalities. The discussions are typically informal, so you may raise your hand to participate in a conversation or simply speak out at appropriate moments.
- 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.**

Students with Disabilities

I would like to hear from anyone who has a disability which may require some modification of seating, testing, or other class requirements so that appropriate arrangements can be made. Please see me after class or during my office hours.

Grievances and Complaints Procedures

If you have any grievance or complaint about course direction, your treatment during class, your assigned marks, or any other problem, please first talk to your professor about the situation. I am very approachable and will work hard to ensure the course is enjoyable for you.
If there is a situation that cannot be resolved in this manner, please contact the Chair of the Department of Computer Science, Professor Jim Cremer, at 14D MacLean Hall, 335-0736.

Academic Dishonesty

A student must submit original work of his or her own construction. Academic dishonesty in the form of copying assignments, projects, or exams from other students or sources is not permitted. If you have any questions about what constitutes academic dishonesty, please contact your professor or consult the printed policy in the *Schedule of Courses* and the *CLAS Bulletin*.

Missing an Exam

Only students who miss an exam for a reason that corresponds to the University of Iowa's policy on "Excused Absences from Examinations" will be permitted to take the exam at a later time. Please note that a make-up exam may have a question format that is different from the regular exam. You must complete an "Explanatory Statement of Absence from Class" form (available at the Registration Center) and present it to the professor for evaluation. If the reason for absence is satisfactory, the student may either take the exam, or if a midterm exam is missed, the midterm exam can be forfeited and 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 covered. The professor is not bound to the topics, timelines, and outline provided as they only serve as a general reference.

Date	Topics Covered and Description
January 20 (T)	**First day of classes. Introduction to course, discuss syllabus/project, WebCT**
January 22 (Th)	Storage issues I: memory hierarchy, hard drive technology, RAID
January 27 (T)	Storage issues II: file organizations, data representation, record types (fixed vs. variable)
January 29 (Th)	Storage issues III: storing records in blocks, file operations, buffering, pointer swizzling
February 3 (T)	Indexing I: motivation, index types, index maintenance, primary/secondary indexes
February 5 (Th)	Indexing II: B-Trees - 2-3 trees, B and B+-Trees
February 10 (T)	Indexing III: Hashing - main memory, external, extendible, linear
February 12 (Th)	Indexing IV: SQL indexing, multi-value indexing (grid files, partitioned hashing) (*Optional topic: indexing spatial data using R-trees*)
February 17 (T)	Query processing I: query processor components, query plans, basic scans, measuring cost of operators, iterators, one-pass algorithms
February 19 (Th)	Query processing II: nested-loop joins, external sorting, two-pass algorithms (sorting) **Project Proposals due.**
February 24 (T)	Query processing III: sort-join, hash-join, two-pass algorithms (hashing), index join
February 26 (Th)	Query optimization I: query parsing, relational algebra laws, heuristic query optimization
March 2 (T)	Query optimization II: physical query plans, pipelining, cost-based query optimization
March 4 (Th)	Transaction processing I: overview, transaction states, ACID properties, schedules
March 9 (T)	Transaction processing II: conflict/view serializablity, precedence graphs **Midterm Exam Review.**
March 11 (Th)	**Midterm Exam In Class.**
March 16 (T)	**No classes during Spring Break.**
March 18 (Th)	**No classes during Spring Break.**
March 23 (T)	Concurrency control I: overview, locks, two-phase locking (2PL), graph protocols
March 25 (Th)	Concurrency control II: multi-granularity locking, timestamps and validation protocols
March 30 (T)	Concurrency control III: deadlock handling, starvation, wait-for graphs
April 1 (Th)	Recovery I: motivation, shadow paging, log-based recovery
April 6 (T)	Recovery II: checkpoints, undo/redo logging, deferred versus immediate update
April 8 (Th)	Introduction to Distributed and Multidatabases
April 13 (T)	*Advanced topic: Concurrency control and recovery in distributed databases (2PC)*
April 15 (Th)	*Advanced topic: Techniques and architectures for data and schema integration*
April 20 (T)	Introduction to Mobile Databases
April 22 (Th)	Project presentations.
April 27 (T)	Project presentations.
April 29 (Th)	Project presentations.
May 4 (T)	Project presentations.
May 6 (Th)	**Final programming project due. Review for final exam. Class evaluations.**
May 13 (Th)	**Final Exam. Thursday, May 13th at 2:15 p.m.**