

Course	ADVANCED DATABASE ORGANIZATION - CS 525 - 01-02-03
Instructor	YOUSEF M ELMEHDWI
Term	SUMMER 2020
Meetings	T/R 6:00-09:10 pm, Online-Blackboard Ultra
Course Webpage	http://cs.iit.edu/~cs525

INSTRUCTOR'S CONTACT INFORMATION	Office Location	Online-Blackboard ✉ yelmehdwi At iit Dot edu
	Office Hours	F 1:00-2:00pm or by appointment

TA'S CONTACT INFORMATION	TBA
-----------------------------	-----

COURSE DESCRIPTION

Databases management systems are a crucial part of most large-scale industry and open-source systems. This course provides comprehensive coverage of issues associated with database system development and an in-depth examination of structures and techniques used in contemporary database management systems (DBMSs). Students will learn about the inner workings of these exciting systems: Which algorithms are used? What are typical architectures used to build a system as complex as a DBMS? What are implementation strategies? These questions and more will be answered during the course.

The course is highly applied, emphasizing practical skills and habits through a series of programming assignments during which students will develop their own tiny DBMS like engine. We will cover the most important aspects/components of a DBMS: **storage and buffer management, indexing, query optimization, query execution, and concurrency control and recovery.**

This course is graduate-level introduction to the design and implementation of data management systems.

ONLINE CLASSES

For CS525, I will live streaming lectures where students are encouraged to attendance classes at the usual course date and time. In addition, all lectures will be recorded and uploaded to course Blackboard right after each class. Students can access the recorded lectures whenever they need them. I prepared and posted a short tutorial on course blackboard under "Assignments & Projects Tutorials" subfolder where I covered the following in details:

- Join Live Sessions and locate recorded lectures materials
- Office hours
- Instructor-to-Student interaction online
- TA-to-Student interaction online
- Participating in discussions

PREREQUISITE(S)

- Courses: CS425
- Programming experience in C, C++ or other low level languages
- Data structures (.e.g., CS401)

COURSE
OBJECTIVES

After attending the course students should:

- Understand the design decisions behind textbook DBMS architectures
- Know the trade-offs of various storage organization techniques
- Be able to build parts of a small-sized data processing system from scratch
- Understand the basics of query optimization
- Know standard implementations of relational operators such as join, aggregation, and set operations
- Be able to estimate the cost of executing an operator/query based on DB statistics
- Know standard database indexing techniques
- Understand concurrency control and recovery mechanisms

SUGGESTED TEXTS,
READINGS, &
MATERIALS

The following text books will be helpful for following the course and studying the presented material. All four textbooks have their merits, but any one should be sufficient as reading material.

- Garcia-Molina, Ullman, and Widom, Database Systems: The Complete Book, 2nd Edition, Prentice Hall, 2008
- Elmasri and Navathe , Fundamentals of Database Systems, 6th Edition , Addison-Wesley , 2003
- Ramakrishnan and Gehrke , Database Management Systems, 3rd Edition , McGraw-Hill , 2002
- Silberschatz, Korth, and Sudarshan , Database System Concepts, 6th Edition , McGraw Hill , 2010

COURSE DETAILS:

Major topics will be covered in the course:

- Disk Storage and Buffer Management
- Indexing and Hashing
- Query Optimization
- Query Execution
- Concurrency control and Recovery
- Advanced Topics (if time permits)

WORKLOAD AND
GRADING POLICIES

Programming Assignments: There will be several programming assignments during the course. Starting from a storage manager you will be implementing your own tiny database-like system from scratch. You will explore how to implement the concepts and data structures discussed in the lectures and readings. The assignments will require the use of skills learned in this course as well as other skills you have developed throughout your program. Each assignment will build upon the code developed during the previous assignment. In the end there will be an optional assignment for extra credit. Each of the regular assignments will have optional parts that give extra credit. All assignments have to be implemented using C/C++.

- **Assignment 1 - Storage Manager:** Implement a storage manager that allows read/writing of blocks to/from a file on disk.
- **Assignment 2 - Buffer Manager:** Implement a buffer manager that manages a buffer of blocks in memory including reading/flushing to disk and block replacement (flushing blocks to disk to make space for reading new blocks from disk).
- **Assignment 3 - Record Manager:** Implement a simple record manager that allows navigation through records, and inserting and deleting records.
- **Assignment 4 - B⁺-Tree Index:** Implement a disk-based B⁺-tree index structure.
 - Implement a disk-based B⁺-tree index structure.

Mid Term and Final Exam: There will be a mid term and a final exam covering the topics of the course.

Grading Policies Each student work product will be graded, and the student's final grade will be determined by assigning each category of work a weighted score according to the following distribution:

Programming Assignments	50% (10% + 10% + 15% + 15%)
Mid Term Exam	20%
Final Exam	30%

LETTER GRADE
DISTRIBUTION

Points	Grade
90 - 100	A
80 - 89	B
70 - 79	C
60 - 69	D
0 - 59	E

OTHER POLICIES

Ethics:

- Any behavior on any assignment or exam that could be considered copying or cheating will result in an immediate zero on the assignment for all parties involved, failure in the class, and notification of the Undergraduate or Graduate Dean's Office. Please see the IIT Code of Academic Honesty <https://web.iit.edu/student-affairs/handbook/fine-print/code-academic-honesty>.
- We will check for plagiarism. Plagiarism will result in zero points for the assignment, potentially academic sanctions, and may result in an E grade.

Late policies

- **-10%** per day late penalty.
- No assignment will be accepted more than 4 calendar days after the original due date!

Make-up Exams

- Only for officially proven health reasons.

Students with Disabilities:

- Reasonable accommodations will be made for students with documented disabilities.
- In order to receive accommodations, students must obtain a letter of accommodation from the Center for Disability Resources.
- The Center for Disability Resources (CDR) is located at 3424 S. State Street - 1C3-2 , 312 567.5744 or disabilities@iit.edu

TENTATIVE COURSE
OUTLINE

The weekly coverage might change as it depends on the progress of the class.

Week	Content
Weeks 1&2	Introduction/ Hardware/File and System Structure
Week 3	Indexing and Hashing
Weeks 4&5	Query Processing
Week 6	Crash Recovery/Concurrency Control

IMPORTANT DATES

Week	Content
05/19	Coding Assignment 1 handed out
05/29	Coding Assignment 1 due
05/29	Coding Assignment 2 handed out
06/07	Coding Assignment 2 due
06/07	Coding Assignment 3 handed out
06/17	Coding Assignment 3 due
06/17	Coding Assignment 4 handed out
06/27	Coding Assignment 4 due
06/12-14	MidTerm Exam, Take home exam
06/25	Final Exam, Take home exam (4 hours)