

Title (Units): **COMP 3220 Database System Implementation (3,3,0)**

Course Aims: To give students a solid background in relational DBMS; to learn general DBMS internals, including data storage, index structures, query evaluation, transaction processing, concurrency control, and crash recovery; and to discuss advanced topics such as distributed databases and data warehouses.

Prerequisite: COMP 1160 Database Management
COMP 1210 Data Structures and Algorithms

Learning Outcomes (LOs):

Upon successful completion of this course, students should be able to:

No.	Learning Outcomes (LOs)
	Knowledge
1	Explain underlying disk-based data storage, organization and access methods
2	Describe query evaluation, query optimization, concurrency control and crash recovery techniques
3	Identify advanced database technologies such as distributed databases and data warehouses
	Professional Skill
4	Analyze the tradeoffs of alternative implementation mechanisms
5	Implement some components of a small relational DBMS
	Attitude
6	Build up team spirit

Calendar Description: This course is to provide an in-depth knowledge of relational database management systems (DBMS). Topics include data storage, index structures, query evaluation, transaction processing, concurrency control, and crash recovery. In addition, advanced topics such as distributed databases and data warehouses will also be covered.

Assessment:

No.	Assessment Methods	Weighting	Remarks
1	Continuous assessment	30%	Continuous assessments are designed to measure how well the students have learned the basic techniques in implementing a database system, including data storage, index structure, query evaluation, transaction processing, crash recovery, etc. A group project is designed to implement some components of a small relational DBMS.
2	Examination	70%	Final examination questions are designed to see how far students have achieved their intended learning outcomes. Questions will primarily be analysis and skills based to assess the students' ability in implementation of database systems.

Rubrics:

Excellent (A)	<ul style="list-style-type: none">Achieve the first five CILOs, demonstrating a good mastery of both the theoretical and practical aspects of database system implementationHave a thorough understanding of concepts and technologies involved in database system implementation, and be able to explain and highlight the key points of these concepts and technologiesAble to draw on a variety of relevant knowledge and concepts and appropriately apply them to analyze the tradeoffs of alternative implementation mechanisms with detailed explanationsAble to implement some components of a small relational database management system according to all requirements
Good (B)	<ul style="list-style-type: none">Achieve the first five CILOs, demonstrating a good understanding of both the theoretical and practical aspects of database system implementationHave a good understanding of concepts and technologies involved in database system

	<p>implementation</p> <ul style="list-style-type: none"> • Able to make use of relevant knowledge and concepts and apply them to analyze the tradeoffs of typical implementation mechanisms with sound explanations • Able to implement some typical components of a small relational database management system according to most of the requirements
Satisfactory (C)	<ul style="list-style-type: none"> • Achieve most of the first five CILOs, demonstrating a basic level of understanding of the theoretical and practical aspects of database system implementation • Have a basic understanding of concepts and technologies involved in database system implementation • Able to make use of some relevant knowledge and concepts and apply them to analyze the tradeoffs of typical implementation mechanisms • Able to implement some typical components of a small relational database management system according to some of the requirements
Marginal Pass (D)	<ul style="list-style-type: none"> • Achieve most of the first five CILOs, with a minimal level of understanding of the theoretical and practical aspects of database system implementation • Have a minimal level of understanding of concepts and technologies involved in database system implementation • Demonstrate the ability to apply some relevant knowledge and concepts to analyze the tradeoffs of a limited number of implementation mechanisms • Able to implement at least one typical component of a small relational database management system according to some minimal requirements
Fail (F)	<ul style="list-style-type: none"> • Achieve less than three of the CILOs, and have little understanding of the theoretical and practical aspects of database system implementation • Unable to provide solutions to simple problems which require basic understanding of concepts and technologies involved in database system implementation • Unable to analyze the tradeoffs of alternative implementation mechanisms • Have little understanding of database implementation technologies and have difficulty in implementing any component of a small relational database management system

Learning Outcomes and Weighting:

Content	LO No.
I. Storage and File Structure	1, 4
II. Indexing and Hashing	1, 4
III. Query Evaluation and Optimization	2, 4,
IV. Transaction Management and Concurrency Control	2, 4
V. Crash Recovery	2, 4
VI. Advanced Topics	3
VII. Group Project	5, 6

References:

Raghu Ramakrishnan and Johannes Gehrke, Database Management Systems, 3rd Edition, McGraw Hill, 2003.

Abraham Silberschatz, Henry F. Korth, and S. Sudarshan, Database System Concepts, 5th Edition, McGraw Hill, 2006.

Ramez Elmasri and Shamkant B. Navathe, Fundamentals of Database Systems, 4th Edition, Addison Wesley, 2003.

C. J. Date, An Introduction to Database Systems, 8th Edition, Addison-Wesley, 2004.

Course Content in Outline:

Topic

- I. Storage and File Structure
 - A. Physical storage media
 - B. Disk space management
 - C. Buffer management
 - D. File organization
 - E. Organization of pages and records in files
- II. Indexing and Hashing
 - A. Indexed sequential access method
 - B. Dynamic tree-structured indexing
 - C. Static hashing
 - D. Dynamic hashing
- III. Query Evaluation and Optimization
 - A. The system catalog
 - B. Query execution plan
 - C. Query operator evaluation
 - D. Query optimization
 - E. External sorting
- IV. Transaction Management and Concurrency Control
 - A. The ACID properties
 - B. Schedule, serializability, and recoverability
 - C. Locking-based concurrency control
 - D. Dealing with deadlocks
 - E. Concurrency control with locking
- V. Crash Recovery
 - A. Type of failures
 - B. The log and other recover-related structures
 - C. Checkpointing
 - D. Recovery techniques
- VI. Advanced Topics
 - A. Database security
 - B. Data mining and data warehousing
 - C. Distributed databases