CS4482: ADVANCED DATABASE SYSTEMS

Effective Term
Semester A 2022/23

Part I Course Overview

Course Title
Advanced Database Systems

Subject Code
CS - Computer Science
Course Number
4482

Academic Unit
Computer Science (CS)

College/School
College of Engineering (EG)

Course Duration
One Semester

Credit Units
3

Level
B1, B2, B3, B4 - Bachelor's Degree

Medium of Instruction
English

Medium of Assessment
English

Prerequisites
CS3402 Database Systems

Precursors
Nil

Equivalent Courses
Nil

Exclusive Courses
Nil

Part II Course Details

Abstract
This course has two objectives. First, it aims to introduce a number of advanced database topics including query optimization and transaction processing. Students are expected to acquire the skill to propose an efficient query evaluation

plan and the knowledge to explain how to enable concurrent transaction executions. Both of them are important to the performance and correctness of a database system. The second objective is to introduce the important issues in various types of specialty databases including parallel databases, distributed databases, spatial databases and real-time databases. We discuss how to exploit parallelism in processing queries, how to process transactions that access data stored at several sites, how to query spatial data and how to process transactions that access temporal data in a timely manner. All these are important techniques to support sophisticated systems or applications like decision support systems, geographic information systems and real-time systems.

Course Intended Learning Outcomes (CILOs)

	CILOs	Weighting (if app.)	DEC-A1	DEC-A2	DEC-A3
1	Develop in-depth understanding of query processing and transaction processing in relational databases.		x		
2	Develop the ability/skill to design efficient evaluation plan to minimize the cost of query execution.			x	
3	Explain the issues that arise in parallel and distributed databases for exploiting parallelism in query processing and enabling distributed transaction processing.			x	
4	Interpret spatial data and temporal data and explain how to perform query and transaction processing on these complex data types.			x	

A1: Attitude
Develop an attitude of discovery/innovation/creativity, as demonstrated by students possessing a strong sense of curiosity, asking questions actively, challenging assumptions or engaging in inquiry together with teachers.

A2: Ability
Develop the ability/skill needed to discover/innovate/create, as demonstrated by students possessing critical thinking skills to assess ideas, acquiring research skills, synthesizing knowledge across disciplines or applying academic knowledge to real-life problems.

A3: Accomplishments
Demonstrate accomplishment of discovery/innovation/creativity through producing /constructing creative works/new artefacts, effective solutions to real-life problems or new processes.

Teaching and Learning Activities (TLAs)

	TLAs	Brief Description	CILO No.	Hours/week (if applicable)
1	Lecture	Explain key concepts.	1, 2, 3, 4	3 hours/week
2	Tutorial	Requires students to individually present and explain his/her solution to a problem.	1, 2, 3, 4	8 hours/semester

3	Individual assignments	Requires students to independently work on 2 assignments. Each assignment contains several questions designed to help students apply the concepts/ algorithms learned to solve practical problems.	1, 2, 3, 4	

Assessment Tasks / Activities (ATs)

	ATs	CILO No.	Weighting (%)	Remarks (e.g. Parameter for GenAI use)
1	Assignments	1, 2, 3	18	
2	Mid-term quiz	1, 2	12	

Continuous Assessment (%)

30

Examination (%)

70

Examination Duration (Hours)

2

Additional Information for ATs

For a student to pass the course, at least 30% of the maximum mark for the examination must be obtained.

Assessment Rubrics (AR)

Assessment Task

Assignments

Criterion

1.1 Ability to design efficient evaluation plan to minimize the cost of query execution
1.2 Ability to apply algorithms to ensure database correctness in the presence of concurrent transaction executions
1.2 Ability to apply algorithms to exploit parallelism in query processing and enable distributed transaction processing in parallel and distributed databases.

Excellent (A+, A, A-)

High

Good (B+, B, B-)

Significant

Fair (C+, C, C-)

Moderate

Marginal (D)

Basic

Failure (F)

Not even reaching marginal levels

Assessment Task

Mid-term Quiz

Criterion

2.1 Ability to apply different algorithms to process an operation
2.2 Ability to evaluate the cost of a query
2.3 Ability to explain how to ensure database correctness in the presence of concurrent transaction executions.

Excellent (A+, A, A-)

High

Good (B+, B, B-)

Significant

Fair (C+, C, C-)

Moderate

Marginal (D)

Basic

Failure (F)

Not even reaching marginal levels

Assessment Task

Examination

Criterion

3.1-3 Same as 2.1-3
3.4 Ability to apply algorithms to exploit parallelism in query processing and enable distributed transaction processing in parallel and distributed databases
3.5 Ability to perform query and transaction processing on spatial data and temporal data.

Excellent (A+, A, A-)

High

Good (B+, B, B-)

Significant

Fair (C+, C, C-)

Moderate

Marginal (D)

Basic

Failure (F)

Not even reaching marginal levels

Part III Other Information

Keyword Syllabus

Query algorithm and optimization. Transaction processing. Distributed database. Parallel database. Spatial database. Real-time database.

Syllabus

- Introduction (review of fundamental database knowledge)
  Relational algebra. Storage. Indexing.
- Query algorithm.
  Selection. External sorting. Join: Block nested loop, Indexed nested loop, Sort merge join, Hash join.
- Query optimization.
  Query cost analysis. Materialization and pipelining. Histogram and size estimation. Cost-based optimization and heuristics.
- Transaction processing.
  ACID properties. Serializability. Concurrency control. Two-phase locking. Deadlock prevention, detection and recovery.
- Distributed database
  Horizontal/vertical partitioning. Commit protocol. Distributed concurrency control. Semi-join.
- Parallel database
  Data partitioning: round robin, hash partition, range partition. Inter-query parallelism. Intra-query parallelism. Intra-operation parallelism. Inter-operation parallelism.
- Spatial database (geographic information system)
  Spatial data. Spatial queries. Indexing: k-d tree, Quadtrees, R-trees.
- Real-time database
  Transaction correctness. Deadlines. Priority assignment. Real-time concurrency control. Temporal consistency. Data-deadline based scheduling.

Reading List

Compulsory Readings

	Title
1	Silberschatz A., Korth H.F. and Sudarshan S. (2011). Database System Concepts. McGraw Hill, 6th edition.

Additional Readings

	Title
1	Hector Garcia-Molina, Jeffrey D. Ullman and Jennifer Widom (2009). Database Systems The Complete Book. Pearson International Edition, 2nd edition.
2	Elmasri R. and Navathe S.B. (2011). Database Systems: Models, Languages, Design and Application Programming. Pearson, 6th edition.