

CS 520: Database Design and Engineering

Objectives

After completing this course, students should be able to:

- Design and model a design scenario using relational data modeling, which includes:
 - Analyze the design anomalies.
 - Analyze and construct functional dependencies, identifying primary keys.
 - Analyze and perform normalization and normal forms.
 - Define referential integrities.
 - Create a normalized relational database.
- Solve abstract relational language, such as relational algebra problems.
- Implement a relational database management system.
- Implement file organizations, indexing, and query processing.
- Implement query optimizations such as rule-based and cost-based.
- Implement integrity, concurrency control, and recovery algorithms.

Prerequisites

- (CS 351 or CS 402 or CS 403) and (CS 430 or CS 406).

Syllabus

- Introduction, basic terminology, and alternate database models: Hierarchical, Network, & OO
- Abstract relational language: relational algebra and relational calculus
- Structured Query Language (SQL)
- Parsing overview - grammars and recursive decent parsers
- File Organization, Indexing, Query Processing: Concept and Implementation
- Relational Data Modeling, keys, referential integrity, constraints, anomalies
- Functional Dependencies, Armstrong axioms, Closure of FD and attribute
- Normalization, Normal Forms, Loss-less decomposition proof, dependency preservation proof
- Query Optimization - cost based; rule based
- Recovery
- Concurrency Control
- Distributed Databases
- Very Large Databases, Data Warehouse, Parallelism
- Paper Presentation
- Project Demo

Edited March 2006 ([html](#), [css](#) checks)