

CS 525: Advanced Database Organization

Objectives

Study of relational, semantic, and object-oriented data models and interfaces. Database management system techniques for query optimization, concurrency control, recovery management, and distributed processing.

Prerequisites

- CS 425.

Syllabus

- Introduction
 - History of database management.
 - Goals of database system development.
- Relational systems
 - Data models.
 - Basic relational concepts.
 - Integrity constraints.
 - Relational algebra.
- Buffer management
 - Goals and importance of buffering.
 - Buffer implementation.
 - Global buffering schemes (FIFO, CLOCK, LFU, LRU, LRU-K).
- Data-store organization
 - Inverted-file organization.
 - Physical representation of attributes.
 - Physical representation of tuples. Internal organization of data pages.
 - Record identification.
 - Mapping short records onto pages.
 - Mapping long records onto pages.
 - Storage allocation and free-space management.
- Indexing and hashing
 - Goals and importance of indexing.
 - Hashing (static hashing, extensible hashing).
 - Comparative indexing schemes (ISAM, B-trees, Prefix B-trees, prefix and trailing key compression).
- Database recovery
 - Concept of a transaction.
 - ACID properties of transactions.
 - Transaction recovery.
 - Physical recovery schemes (Redo/No-undo, Undo/No-Redo, Undo/Redo, No-Undo/No-Redo).
- Concurrency control
 - Transactions and concurrency.
 - Schedules and serializability.
 - Bad dependencies.
 - Locking (lock modes, lock duration, lock granularity, two-phase locking, isolation levels).
 - Lock implementation (hard locks, intention locks, lock queues, lock management and structures).
- Page-level logical recovery
 - Brief overview.
 - Basic structures.
 - Recovery operations (forward processing, rollback, checkpointing, and restart processing).
- Query Processing
- Query Optimization
- Distributed databases
 - Date's requirements for distributed data management.
 - Problems of distributed database management
 - Object naming.
 - Data-dictionary management.
 - Data fragmentation.
 - Distributed query processing and optimization.
 - Global transactions and two-phase commit.
- Data replication
 - Objectives and requirements of data replication.
 - Replication schemes.
 - Synchronous replication.

- Periodic state-based replication.
 - Asynchronous replication.
 - Symmetric replication.
 - Evaluation of different replication schemes.
- System Architecture
 - Client/server architecture.
 - Processes and multi-process server architectures.
 - Threading and multi-threaded server architectures.
- Research Topics in Database Systems

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