

Lectures for the course: Advanced Database Systems (IT 60113)

Week 1

Lecture 1 – 21/07/2010

- Introduction to the course
- Expectations
- Evaluation Guideline

Lecture 2 – 22/07/2010

- Introduction to ER Model
- Entity Sets, Relationship Sets, Attributes
- Multivalued Attributes, Derived Attributes, Composite Attributes
- Cardinality, Participation
- Keys

Week 2

Lecture 3 – 26/07/2010

- Keys of relationship sets
- Design issues
- Weak entities
- E-R Diagram
- Ternary relationships

Lecture 4+5 – 28/07/2010

- Reflexive relationships
- Generalization and specialization
- ER Model exercise

Lecture 6 – 29/07/2010

- Mapping ER Model to Tables
- Redundancy of Tables
- Alternative representation of tables mapped from generalization and specialization

Week 3

Lecture 7 – 02/08/2010

- Relational model
- Database schema
- Foreign keys
- Relational algebra

Lecture 8+9 – 03/08/2010

- Basics of SQL
- Simple Select
- Joining of two or more tables
- Aliasing
- Aggregation
- Group By

Lecture 10 – 04/08/2010

- Introduction to Active databases
- STARBURST Rules
- Syntax and Semantics
- Example STARBURST rules
- Rule prioritization

Week 4

Lecture 11 – 09/08/2010

- STARBURST Semantics continued
- Net Effect
- INSERTED, DELETED, OLD_UPDATED and NEW_UPDATED tables
- Multiple rules

Lecture 12+13 – 11/08/2010

- Oracle Triggers
- Syntax and Semantics of Oracle Triggers
- Examples
- Active Rule Design - Termination

Lecture 14 – 12/08/2010

- Confluence
- Observable Determinism
- Applications of Active Rules
- Integrity maintenance
- Derived data maintenance
- Replication

Week 5

Lecture 15 – 16/08/2010

- Introduction to Temporal databases
- Notion of time and basic SQL support for handling time
- Examples on processing time related data in standard SQL

Lecture 16+17 – 18/08/2010

- Class Test 1 held

Lecture 18 – 19/08/2010

- Further examples on processing time related data in standard SQL
- Class test 1 scripts shown

Week 6

Lecture 19 – 23/08/2010

- Time as a database dimension and its interpretation
- Chronon, anchored and un-anchored time
- Valid time and transaction time
- Maintaining history with valid time and transaction time

Lecture 20+21 – 25/08/2010

- Bitemporal Conceptual Data Model
- Other temporal data models
- Value equivalent tuples
- Introduction to TSQL2
- CREATE table and basic select queries

Lecture 22 – 26/08/2010

- Simple TSQL2 syntax and semantics
- Coalescing at the time of returning query results

- Restructuring
- Partitioning

Week 7

Lecture 23 – 30/08/2010

- VALID clause to specify valid time of tuples in Select query
- Insert, Delete and update
- Effect of update resulting in insert

Lecture 24+25 – 01/09/2010

- Aggregation and group by
- Transaction time support
- Valid event relations
- Review of term paper status

Lecture 26 – 02/09/2010

- Introduction to datalog
- Facts, base predicates and derived predicates
- Rules – Head and Body of rules, Goals
- Conjunction and Disjunction
- Use of wildcard

Week 8

Lecture 27 – 06/09/2010

- More examples of derived predicates
- Negation
- Use of negation for universal quantification – All

Lecture 28+29 – 08/09/2010

- Recursive queries
- Bill of Materials example
- Safe Datalog programs

Lecture 30 – 09/09/2010

- Mapping safe non-recursive datalog programs to relational algebra
- Mapping datalog programs to SQL
- Summary of topics covered up to mid-sem exam

- Clarification to queries on topics covered before mid-sem exam

Week 9

13/09/2010 – 20/09/2010 Mid-sem Exam

Week 10

Lecture 31+32 – 22/09/2010

- Introduction to transactions
- Properties of transactions
- Isolation and concurrency control
- Schedules
- Serial schedules
- Concurrent schedule
- Exam scripts shown and feedback given

Lecture 33 – 23/09/2010

- More examples of concurrent schedule
- Conflict equivalent schedule
- Conflict serializability
- Feedback on term paper given

Week 11

Lecture 34 – 27/09/2010

- View serializability
- Testing for conflict serializability and view serializability
- Recoverable and cascadeless schedules

Lecture 35+36 – 29/09/2010

- Concurrency Control
- Lock based protocols
- Shared and exclusive locks
- Deadlock
- Two phase locking
- Starvation
- S2PL and R2PL
- Lock upgrading

Lecture 37 – 30/09/2010

- Tree based locking protocol
- Deadlock free protocol

Week 12

Lecture 38 – 04/10/2010

- Timestamp based protocol
- Validation based protocol
- Optimistic concurrency control

Lecture 39+40 – 06/10/2010

- Multi-level concurrency control
- Multi-version concurrency control

Lecture 41 – 07/10/2010

- Effect of insert and delete on concurrency control
- Phantom phenomenon
- Index locking scheme

Lecture 42 – 07/10/2010 (Compensatory Lecture)

- Weak Levels of consistency
- Type 2 consistency
- Cursor stability
- Repeatable read
- Read committed
- Read uncommitted

Week 13

Lecture 43 – 11/10/2010

- Introduction to distributed databases
- Replication and fragmentation
- Relative advantages and disadvantages
- Real-life examples of distributed databases

Week 14

Lecture 44+45 – 20/10/2010

- Recap of distributed database architecture
- Two phase commit and three phase commit
- Type of failure
- Effect of failure on two phase and three phase commit protocols
- Writing lock information in log
- Concurrency control in distributed databases
- Centralized and distributed locking
- Primary copy, majority protocol and biased protocol

Lecture 46 – 21/10/2010

- Distributed queries
- Distributed join processing
- Semi join processing

Week 15

Lecture 47 – 25/10/2010

- Indexing multidimensional data
- Nature of multidimensional data
- Multidimensional point data
- Nature of queries in multidimensional data
- Shortcomings of B/B+ tree in handling multidimensional data
- Introduction to R-Tree

Lecture 48+49 – 27/10/2010

- Class Test 2 held

Lecture 50 – 28/10/2010

- R-Tree node structure
- Formal definition of R-tree
- Class Test 2 scripts shown

Week 16

Lecture 51 – 01/11/2010

- Searching on R-Tree
- Insert in R-Tree
- ChooseLeaf and AdjustTree

- SplitNode – exhaustive search

Lecture 52+53 – 03/11/2010

- SplitNode – Quadratic cost algorithm
- Example of building R-tree
- Nearest neighbor search in R-Tree
- MinDist between a query point and a hyper rectangle
- Idea of nearest neighbor hyper sphere

Lecture 54 – 03/11/2010

- MinMaxDist between a query point and a hyper rectangle
- Distance between a query point and an object
- Search ordering by MinDist and MinMaxDist and their impact
- Different pruning strategies

Week 17

Lecture 55 – 08/11/2010

- K-NN search
- Variants of R-Tree – R*-Tree
- Other distance measures like cosine distance
- Multimedia databases – R-Tree for high dimensional point data

Lecture 56+57 – 10/11/2010

- SS-Tree
- Properties of high dimensional data
- Approximate search techniques
- Directions towards indexing data in NDDS and Hybrid data space
- Term project presentations

Lecture 58 – 11/11/2010

- Term project presentations

Week 18

Lecture 59 – 15/11/2010

- Summary of topics covered