INTRODUCTORY VERSION

Database Systems
An Application-Oriented Approach
SECOND EDITION

Chapter 2

The Big Picture

Michael KIFER   Arthur BERNSTEIN   Philip M. LEWIS

PEARSON
Addison
Wesley

Databases

- We are particularly interested in relational databases
- Data is stored in tables.

Table

- Set of rows (no duplicates)
- Each *row* describes a different entity
- Each *column* states a particular fact about each entity
  - Each column has an associated *domain*

Id	Name	Address	Status
1111	John	123 Main	fresh
2222	Mary	321 Oak	soph
1234	Bob	444 Pine	soph
9999	Joan	777 Grand	senior

- Domain of *Status* = {fresh, soph, junior, senior}

Relation

- Mathematical entity corresponding to a table
  - row ~ *tuple*
  - column ~ *attribute*
- Values in a tuple are *related* to each other
  - John lives at 123 Main
- Relation **R** can be thought of as predicate **R**
  - **R**(x,y,z) is true iff tuple (x,y,z) is in **R**

Operations

- Operations on relations are precisely defined
  - Take relation(s) as argument, produce new relation as result
  - Unary (*e.g.*, delete certain rows)
  - Binary (*e.g.*, union, Cartesian product)
- Corresponding operations defined on tables as well
- Using mathematical properties, *equivalence* can be decided
  - Important for *query optimization:*

$$op1(T1,op2(T2)) \overset{?}{=} op3(op2(T1),T2)$$

Structured Query Language: SQL

- Language for manipulating tables
- *Declarative* – Statement specifies *what* needs to be obtained, *not how* it is to be achieved (*e.g.*, how to access data, the order of operations)
- Due to declarativity of SQL, DBMS determines evaluation strategy
  - This greatly simplifies application programs
  - *But DBMS is not infallible*: programmers should have an idea of strategies used by DBMS so they can design better tables, indices, statements, in such a way that DBMS can evaluate statements efficiently

Structured Query Language (SQL)

SELECT <*attribute list*>
FROM <*table list* >
WHERE <*condition*>

- Language for constructing a new table from argument table(s).
  - FROM indicates source tables
  - WHERE indicates which *rows* to retain
    - It acts as a filter
  - SELECT indicates which *columns* to extract from retained rows
    - Projection
- The result is a table.

2-7

Example

SELECT *Name*
FROM *Student*
WHERE *Id* > 4999

Id	Name	Address	Status
1234	John	123 Main	fresh
5522	Mary	77 Pine	senior
9876	Bill	83 Oak	junior

Student

Name
Mary
Bill

Result

2-8

Examples

SELECT *Id*, *Name* FROM Student

SELECT *Id*, *Name* FROM Student
  WHERE *Status* = 'senior'

SELECT * FROM Student
  WHERE *Status* = 'senior'

result is a table with one column and one row

SELECT COUNT(*) FROM Student
  WHERE *Status* = 'senior'

2-9

More Complex Example

- **Goal:** table in which each row names a senior and gives a course taken and grade
- Combines information in two tables:
  - Student: *Id*, *Name*, *Address*, *Status*
  - Transcript: *StudId*, *CrsCode*, *Semester*, *Grade*

SELECT *Name*, *CrsCode*, *Grade*
FROM Student, Transcript
WHERE *StudId* = *Id* AND *Status* = '*senior*'

2-10

Join

SELECT *a1*, *b1*
FROM T1, T2
WHERE *a2* = *b2*

T1

a1	a2	a3
A	1	xxy
B	17	rst

T2

b1	b2
3.2	17
4.8	17

FROM T1, T2
yields:

a1	a2	a3	b1	b2
A	1	xxy	3.2	17
A	1	xxy	4.8	17
B	17	rst	3.2	17
B	17	rst	4.8	17

WHERE *a2* = *b2*
yields:

a1	a2	a3	b1	b2
B	17	rst	3.2	17
B	17	rst	4.8	17

SELECT *a1*, *b1*
yields result:

B	3.2
B	4.8

2-11

Modifying Tables

UPDATE *Student*
SET *Status* = 'soph'
WHERE *Id* = 111111111

INSERT INTO Student (*Id*, *Name*, *Address*, *Status*)
VALUES (999999999, 'Bill', '432 Pine', 'senior')

DELETE FROM Student
WHERE *Id* = 111111111

2-12

Creating Tables

CREATE TABLE Student (
  Id INTEGER,
  Name CHAR(20),
  Address CHAR(50),
  Status CHAR(10),
  PRIMARY KEY (*Id*))

Constraint: explained later

2-13

Transactions

- Many enterprises use databases to store information about their state
  - *E.g.*, balances of all depositors
- The occurrence of a real-world event that changes the enterprise state requires the execution of a program that changes the database state in a corresponding way
  - *E.g.*, balance must be updated when you deposit
- A *transaction* is a program that accesses the database in response to real-world events

2-14

Transactions

- Transactions are <u>not just ordinary programs</u>
- Additional requirements are placed on transactions (and particularly their execution environment) that go beyond the requirements placed on ordinary programs.
  - **A**tomicity
  - **C**onsistency
  - **I**solation   *ACID properties*
  - **D**urability
  (explained next)

2-15

Integrity Constraints

- Rules of the enterprise generally limit the occurrence of certain real-world events.
  - Student cannot register for a course if current number of registrants = maximum allowed
- Correspondingly, allowable database states are restricted.
  - *cur_reg <= max_reg*
- These limitations are expressed as *integrity constraints,* which are assertions that must be satisfied by the database state.

2-16

Consistency

- Transaction designer must ensure that
  IF the database is in a state that satisfies all integrity constraints when execution of a transaction is started
  THEN when the transaction completes:
    - All integrity constraints are once again satisfied (constraints can be violated in intermediate states)
    - New database state satisfies specifications of transaction

2-17

Atomicity

- A real-world event either happens or does not happen.
  - Student either registers or does not register.
- Similarly, the system must ensure that either the transaction runs to completion (*commits*) or, if it does not complete, it has no effect at all (*aborts*).
  - This is not true of ordinary programs. A hardware or software failure could leave files partially updated.

2-18

Durability

- The system must ensure that once a transaction commits its effect on the database state is not lost in spite of subsequent failures.
  - Not true of ordinary systems. For example, a media failure after a program terminates could cause the file system to be restored to a state that preceded the execution of the program.

2-19

Isolation

- Deals with the execution of multiple transactions concurrently.
- If the initial database state is consistent and accurately reflects the real-world state, then the *serial* (one after another) execution of a set of consistent transactions preserves consistency.
- But serial execution is *inadequate* from a performance perspective.

2-20

Concurrent Transaction Execution
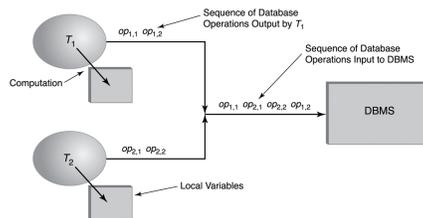
2-21

Isolation

- Concurrent (interleaved) execution of a set of transactions offers performance benefits, but might not be correct.
- **Example**: Two students execute the course registration transaction at about the same time

  (*cur_reg* is the number of current registrants)

T_1: read(*cur_reg* : 29)   write(*cur_reg* : 30)

T_2:   read(*cur_reg* : 29) write(*cur_reg* : 30)

  time $\rightarrow$

Result: Database state no longer corresponds to real-world state, integrity constraint violated.

2-22

Isolation

- The effect of concurrently executing a set of transactions must be the same as if they had executed serially in some order
  - The execution is thus *not* serial, but *serializable*
- Serializable execution has better performance than serial, but performance might still be inadequate. Database systems offer several isolation levels with different performance characteristics (but some guarantee correctness only for certain kinds of transactions – not in general)

2-23

ACID Properties

- The transaction monitor is responsible for ensuring atomicity, durability, and (the requested level of) isolation.
  - Hence it provides the abstraction of failure-free, non-concurrent environment, greatly simplifying the task of the transaction designer.
- The transaction designer is responsible for ensuring the consistency of each transaction, but doesn't need to worry about concurrency and system failures.

2-24