CS425 – Fall 2014
Boris Glavic
Course Information

Hi, I am **Boris Glavic, Assistant Professor** in **CS**

0.4

©Silberschatz, Korth and Sudarshan

Why are Databases Important?

- **What do Databases do?**

1. Provide persistent storage

2. Efficient declarative access to data -> Querying

3. Protection from hardware/software failures

4. Safe concurrent access to data

Who uses Databases?

- **Most big software systems involve DBs!**
  - Business Intelligence ⇒ e.g., IBM Cognos
  - Web based systems
  - …
- **You!** (desktop software)
  - Your music player ⇒ e.g., Amarok
  - Your Web Content Management System
  - Your email client
  - …
- Every big company
  - Banks
  - Insurance
  - Government
  - Google, …
  - …

Who Produces Databases?

- **Traditional relational database systems is big business**
  - IBM ⇒ DB2
  - Oracle ⇒ Oracle ☺
  - Microsoft ⇒ SQLServer
  - Open Source ⇒ MySQL, Postgres, …
- **Emerging distributed systems with DB characteristics and Big Data**
  - Cloud storage and Key-value stores ⇒Amazon S3, Google Big Table, . . .
  - Big Data Analytics ⇒Hadoop, Google Map & Reduce, . . .
  - SQL over Distributed Platforms ⇒ Hive, Tenzing, …

Why are Database Interesting (for Students)?

- **The pragmatic perspective**
  - Background in databases make you competitive in the job market ;-)

- **Systems and theoretical research**
  - Database research has a strong systems aspect
    - Hacking complex and large systems
    - Low-level optimization
      - cache-conscious algorithms
      - Exploit modern hardware
  - Databases have a strong theoretical foundation
    - Complexity of query answering
    - Expressiveness of query languages
    - Concurrency theory
    - …

Why are Database Interesting (for Students)?

- Connection to many CS fields
  - Distributed systems
    - Getting more and more important
  - Compilers
  - Modeling
  - AI and machine learning
    - Data mining
  - Operating and file systems
  - Hardware
    - Hardware-software co-design

Webpage and Faculty

- **Course Info**
  - **Course Webpage**: http://cs.iit.edu/~cs425
  - **Google Group**:
    https://groups.google.com/d/forum/cs425-2014-fall-group
    - ▸ Used for announcements
    - ▸ Use it to discuss with me, TA, and fellow students
  - **Syllabus:** http://cs.iit.edu/~cs425/files/syllabus.pdf
- **Faculty**
  - **Boris Glavic** (http://cs.iit.edu/~glavic)
  - **Email:** bglavic@iit.edu
  - **Phone**: 312.567.5205
  - **Office**: Stuart Building, room 226C
  - **Office Hours**: Tuesdays, 12pm-1pm (and by appointment)

TAs

- **Tas**
  - **TBA**

Workload and Grading

- **Exams**
  - Midterm (25%)
  - Final (35%)

- **Homework Assignments** (preparation for exams!)
  - HW1 (Relational algebra)
  - HW2 (SQL)
  - HW3 (Database modeling)

- **Course Project**
  - In groups of 3 students
  - Given an example application (e.g., ticketing system)
    - Develop a database model
    - Derive a database schema from the model
    - Implement the application accessing the database

Course Objectives

- Understand the underlying ideas of database systems

- Understand the **relational data model**

- Be able to write and understand **SQL** queries and data definition statements

- Understand **relational algebra** and its connection to SQL

- Understand how to **write programs that access a database server**

- Understand the **ER model** used in database design

- Understand **normalization** of database schemata

- Be able to **create a database design** from a requirement analysis for a specific domain

- Know basic **index structures** and understand their importance

- Have a basic understanding of relational database concepts such as **concurrency control**, **recovery**, **query processing**, and **access control**

Course Project

- Forming groups
  - Your responsibility!
  - Inform me + TA
  - Deadline: Sep 8th
- Oracle Server Accounts
- Git repositories
  - Create an account on Bitbucket.org (https://bitbucket.org/)
  - We will create a repository for each student
  - Use it to exchange code with your fellow group members
  - The project has to be submitted via the group repository
- Timeline:
  - Brainstorming on application (by Sep 11th)
  - Design database model (by Nov 12th)
  - Derive relational model (by Nov 25th)
  - Implement application (by end of the semester)

Fraud and Late Assignments

- All work has to be original!
  - Cheating = 0 points for assignment/exam
  - Possibly E in course and further administrative sanctions
  - Every dishonesty will be reported to office of academic honesty
- Late policy:
  - -20% per day
  - No exceptions!
- Course projects:
  - Every student has to contribute in **every** phase of the project!
  - **Don't let others freeload on you hard work!**
    - ‣ Inform me or TA immediatly

Reading and Prerequisites

- **Textbook:** Silberschatz, Korth and Sudarsham
  - ***Database System Concepts, 6th edition***
  - McGraw Hill
  - publication date:2006,
  - ISBN 0-13-0-13-142938-8.
- Prerequisites:
  - CS 331 or CS401 or CS403

Outline

- Introduction
- Relational Data Model
- Formal Relational Languages (relational algebra)
- SQL
- Database Design
- Transaction Processing, Recovery, and Concurrency Control
- Storage and File Structures
- Indexing and Hashing
- Query Processing and Optimization

CS425 – Fall 2014
Boris Glavic
Chapter 1: Introduction

Modified from:

Database System Concepts, 6th Ed.

Textbook: Chapter 1

Database Management System (DBMS)

■ DBMS contains information about a particular domain

- Collection of interrelated data
- Set of programs to access the data
- An environment that is both *convenient* and *efficient* to use

■ Database Applications:

- Banking: transactions
- Airlines: reservations, schedules
- Universities: registration, grades
- Sales: customers, products, purchases
- Online retailers: order tracking, customized recommendations
- Manufacturing: production, inventory, orders, supply chain
- Human resources: employee records, salaries, tax deductions

■ Databases can be very large.

■ Databases touch all aspects of our lives

University Database Example

- Application program examples

  - Add new students, instructors, and courses

  - Register students for courses, and generate class rosters

  - Assign grades to students, compute grade point averages (GPA) and generate transcripts

- In the early days, database applications were built directly on top of file systems

Drawbacks of using file systems to store data

- Data redundancy and inconsistency
    - Multiple file formats, duplication of information in different files
- Difficulty in accessing data
    - Need to write a new program to carry out each new task
- Data isolation — multiple files and formats
- Integrity problems
    - Integrity constraints (e.g., account balance > 0) become "buried" in program code rather than being stated explicitly
    - Hard to add new constraints or change existing ones

Drawbacks of using file systems to store data (Cont.)

- Atomicity of updates
  - ▸ Failures may leave database in an inconsistent state with partial updates carried out
  - ▸ Example: Transfer of funds from one account to another should either complete or not happen at all

- Concurrent access by multiple users
  - ▸ Concurrent access needed for performance
  - ▸ Uncontrolled concurrent accesses can lead to inconsistencies
    - – Example: Two people reading a balance (say 100) and updating it by withdrawing money (say 50 each) at the same time

- Security problems
  - ▸ Hard to provide user access to some, but not all, data

Database systems offer solutions to all the above problems!

Levels of Abstraction

- **Physical level:** describes how a record (e.g., customer) is stored.

- **Logical level:** describes data stored in database, and the relationships among the data.

    type *instructor* = **record**

    > *ID* : string;
    > *name* : string;
   > *dept_name* : string;
    > *salary* : integer;

    end;

- **View level:** application programs hide details of data types. Views can also hide information (such as an employee's salary) for security purposes.

View of Data

An architecture for a database system

Instances and Schemas

- Similar to types and variables in programming languages

- **Schema** – the logical structure of the database
  - Example: The database consists of information about a set of customers and accounts and the relationship between them
  - Analogous to type information of a variable in a program
  - **Physical schema**: database design at the physical level
  - **Logical schema**: database design at the logical level

- **Instance** – the actual content of the database at a particular point in time
  - Analogous to the value of a variable

- **Physical Data Independence** – the ability to modify the physical schema without changing the logical schema
  - Applications depend on the logical schema
  - In general, the interfaces between the various levels and components should be well defined so that changes in some parts do not seriously influence others.

- **Logical Data Independence** – the ability to modify the logical schema without changing the applications
  - For example, add new information to each employee

Data Models

- A collection of tools for describing
  - Data
  - Data relationships
  - Data semantics
  - Data constraints
- Relational model
- Entity-Relationship data model (mainly for database design)
- Object-based data models (Object-oriented and Object-relational)
- Semistructured data model (XML)
- Other older models:
  - Network model
  - Hierarchical model
- Other newer (or revived) models:
  - Key-value

Relational Model

- Relational model (Chapter 2)
- Example of tabular data in the relational model

Columns (attributes)

ID	name	dept_name	salary
22222	Einstein	Physics	95000
12121	Wu	Finance	90000
32343	El Said	History	60000
45565	Katz	Comp. Sci.	75000
98345	Kim	Elec. Eng.	80000
76766	Crick	Biology	72000
10101	Srinivasan	Comp. Sci.	65000
58583	Califieri	History	62000
83821	Brandt	Comp. Sci.	92000
15151	Mozart	Music	40000
33456	Gold	Physics	87000
76543	Singh	Finance	80000

Rows (tuples)

(a) The *instructor* table

A Sample Relational Database

ID	name	dept_name	salary
22222	Einstein	Physics	95000
12121	Wu	Finance	90000
32343	El Said	History	60000
45565	Katz	Comp. Sci.	75000
98345	Kim	Elec. Eng.	80000
76766	Crick	Biology	72000
10101	Srinivasan	Comp. Sci.	65000
58583	Califieri	History	62000
83821	Brandt	Comp. Sci.	92000
15151	Mozart	Music	40000
33456	Gold	Physics	87000
76543	Singh	Finance	80000

(a) The *instructor* table

dept_name	building	budget
Comp. Sci.	Taylor	100000
Biology	Watson	90000
Elec. Eng.	Taylor	85000
Music	Packard	80000
Finance	Painter	120000
History	Painter	50000
Physics	Watson	70000

(b) The *department* table

Data Manipulation Language (DML)

- Language for accessing and manipulating the data organized by the appropriate data model

  - DML also known as query language

- Two classes of languages

  - **Procedural** – user specifies what data is required and how to get those data

  - **Declarative (nonprocedural)** – user specifies what data is required without specifying how to get those data

- **SQL** is the most widely used query language

Data Definition Language (DDL)

- Specification notation for defining the database schema

  Example: **create table** *instructor* (
                 ID   **char**(5),
                 name   **varchar**(20)**,**
                 dept_name **varchar**(20),
                 salary   **numeric**(8,2))

- DDL compiler generates a set of table templates stored in a ***data dictionary***

- Data dictionary contains **metadata** (i.e., data about data)

  - Database schema

  - **Integrity constraints**

    ▸ Primary key (ID uniquely identifies instructors)

    ▸ Referential integrity (**references** constraint in SQL)

       – e.g. *dept_name* value in any *instructor* tuple must appear in *department* relation

  - Authorization

SQL

- **SQL**: widely used **declarative** (non-procedural) language
  - Example: Find the name of the instructor with ID 22222
    select   *name*
    from   *instructor*
    where   *instructor.ID* = '22222'
  - Example: Find the ID and building of instructors in the Physics dept.

    select *instructor.ID*, *department.building*
    from *instructor*, *department*
    where *instructor.dept_name = department.dept_name* **and**
             department.dept_name = 'Physics'

- Application programs generally access databases through one of
  - Language extensions to allow embedded SQL
  - Application program interface (e.g., ODBC/JDBC) which allow SQL queries to be sent to a database
- Chapters 3, 4 and 5

Database Design

The process of designing the general structure of a database:

■ Logical Design – Deciding on the database schema. Database design requires that we find a "good" representation of the information from an application domain (e.g., banking) as a collection of relation schemas.

- ● Business decision – What information should we record in the database?

- ● Computer Science decision – What relation schemas should we have and how should the attributes be distributed among the various relation schemas?

■ Physical Design – Deciding on the physical layout of the database

Database Design?

- Is there any problem with this design?

ID	name	salary	dept_name	building	budget
22222	Einstein	95000	Physics	Watson	70000
12121	Wu	90000	Finance	Painter	120000
32343	El Said	60000	History	Painter	50000
45565	Katz	75000	Comp. Sci.	Taylor	100000
98345	Kim	80000	Elec. Eng.	Taylor	85000
76766	Crick	72000	Biology	Watson	90000
10101	Srinivasan	65000	Comp. Sci.	Taylor	100000
58583	Califieri	62000	History	Painter	50000
83821	Brandt	92000	Comp. Sci	Taylor	100000
15151	Mozart	40000	Music	Packard	80000
33456	Gold	87000	Physics	Watson	70000
76543	Singh	80000	Finance	Painter	120000

Database Design?

- **Example:** Changing the budget of the 'Physics' department

  - Updates to many rows!

  - Easy to break **integrity**

    - If we forget to update a row, then we have multiple budget values for the physics department!

| ID	name	salary	dept_name	building	budget
22222	Einstein	95000	Physics	Watson	70000
12121	Wu	90000	Finance	Painter	120000
32343	El Said	60000	History	Painter	50000
45565	Katz	75000	Comp. Sci.	Taylor	100000
98345	Kim	80000	Elec. Eng.	Taylor	85000
76766	Crick	72000	Biology	Watson	90000
10101	Srinivasan	65000	Comp. Sci.	Taylor	100000
58583	Califieri	62000	History	Painter	50000
83821	Brandt	92000	Comp. Sci	Taylor	100000
15151	Mozart	40000	Music	Packard	80000
33456	Gold	87000	Physics	Watson	70000
76543	Singh	80000	Finance	Painter	120000

Design Approaches

■ Normalization Theory (Chapter 8)

   ● Formalize what designs are bad, and test for them

■ Entity Relationship Model (Chapter 7)

   ● Models an enterprise as a collection of *entities* and *relationships*

      ▸ Entity: a "thing" or "object" in the enterprise that is distinguishable from other objects

        – Described by a set of *attributes*

      ▸ Relationship: an association among several entities

   ● Represented diagrammatically by an *entity-relationship diagram:*

The Entity-Relationship Model

- Models an enterprise as a collection of *entities* and *relationships*
  - Entity: a "thing" or "object" in the enterprise that is distinguishable from other objects
    - Described by a set of *attributes*
  - Relationship: an association among several entities
- Represented diagrammatically by an *entity-relationship diagram:*

instructor
<u>ID</u> name salary

member

department
<u>dept_name</u> building budget

What happened to dept_name of instructor and student?

Object-Relational Data Models

■ Relational model: flat, "atomic" values

● E.g., integer

■ Object Relational Data Models

● Extend the relational data model by including object orientation and constructs to deal with added data types.

● Allow attributes of tuples to have complex types, including non-atomic values such as nested relations.

● Preserve relational foundations, in particular the declarative access to data, while extending modeling power.

● Provide upward compatibility with existing relational languages.

XML: Extensible Markup Language

- Defined by the WWW Consortium (W3C)

- Originally intended as a document markup language not a database language

- The ability to specify new tags, and to create nested tag structures made XML a great way to exchange **data**, not just documents

- XML has become the basis for all new generation data interchange formats.

- A wide variety of tools is available for parsing, browsing and querying XML documents/data

Storage Management

- **Storage manager** is a program module that provides the interface between the low-level data stored in the database (on disk) and the application programs and queries submitted to the system.

- The storage manager is responsible to the following tasks:
  - Interaction with the file manager
  - Efficient storing, retrieving and updating of data

- Issues:
  - Storage access
  - File organization
  - Indexing and hashing

Query Processing

1. Parsing and translation
2. Optimization
3. Evaluation
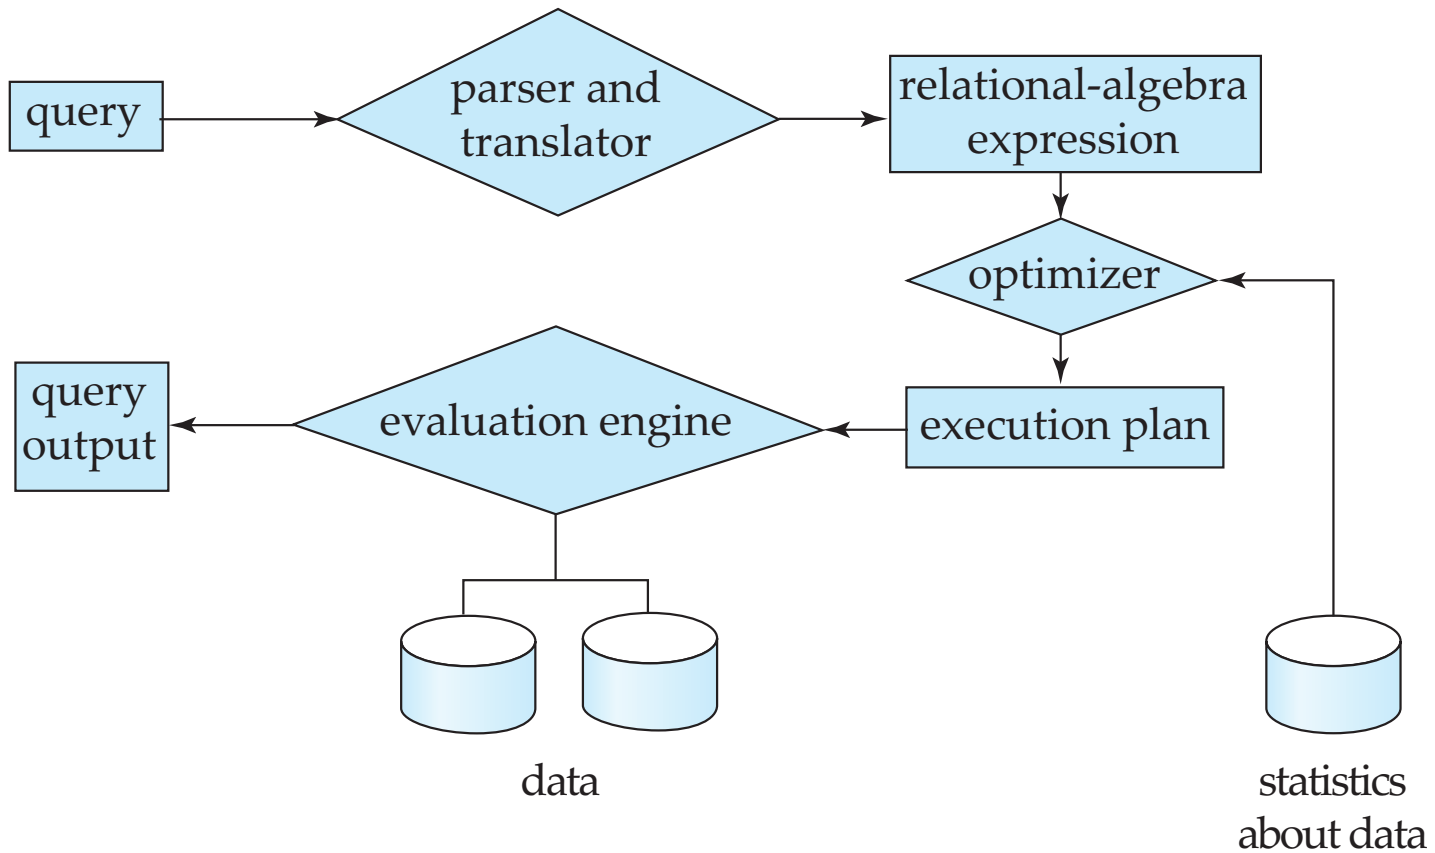
Query Processing (Cont.)

- Alternative ways of evaluating a given query
  - Equivalent expressions
  - Different algorithms for each operation
- Cost difference between a good and a bad way of evaluating a query can be enormous
- Need to estimate the cost of operations
  - Depends critically on statistical information about relations which the database must maintain
  - Need to estimate statistics for intermediate results to compute cost of complex expressions
- Need to search for a good plan (low costs)
  - Traversing the search space of alternative ways (plans) to compute the query result
 - This is called **query optimization**

Transaction Management

- What if the system fails?

- What if more than one user is concurrently updating the same data?

- A **transaction** is a collection of operations that performs a single logical function in a database application

- **Transaction-management component** ensures that the database remains in a consistent (correct) state despite system failures (e.g., power failures and operating system crashes) and transaction failures.

- **Concurrency-control manager** controls the interaction among the concurrent transactions, to ensure the consistency of the database.
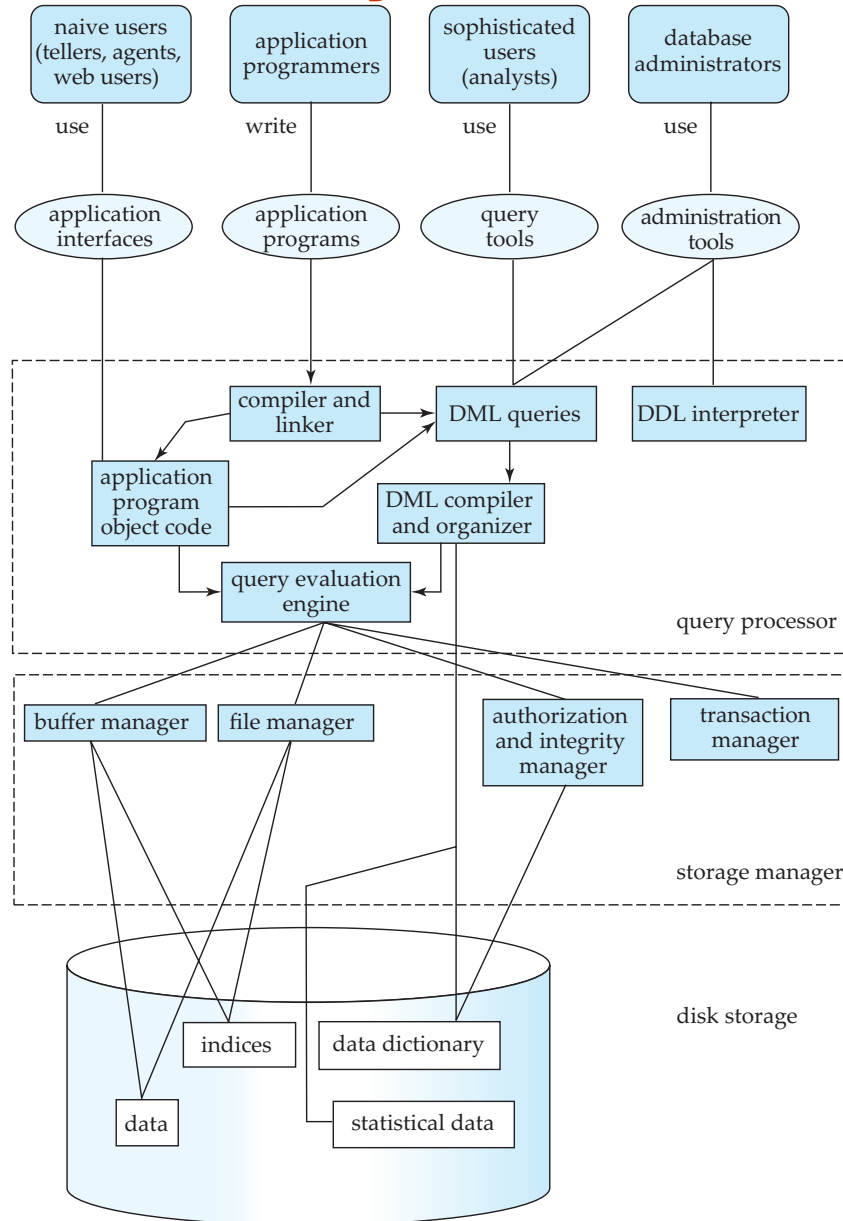
Database Users and Administrators
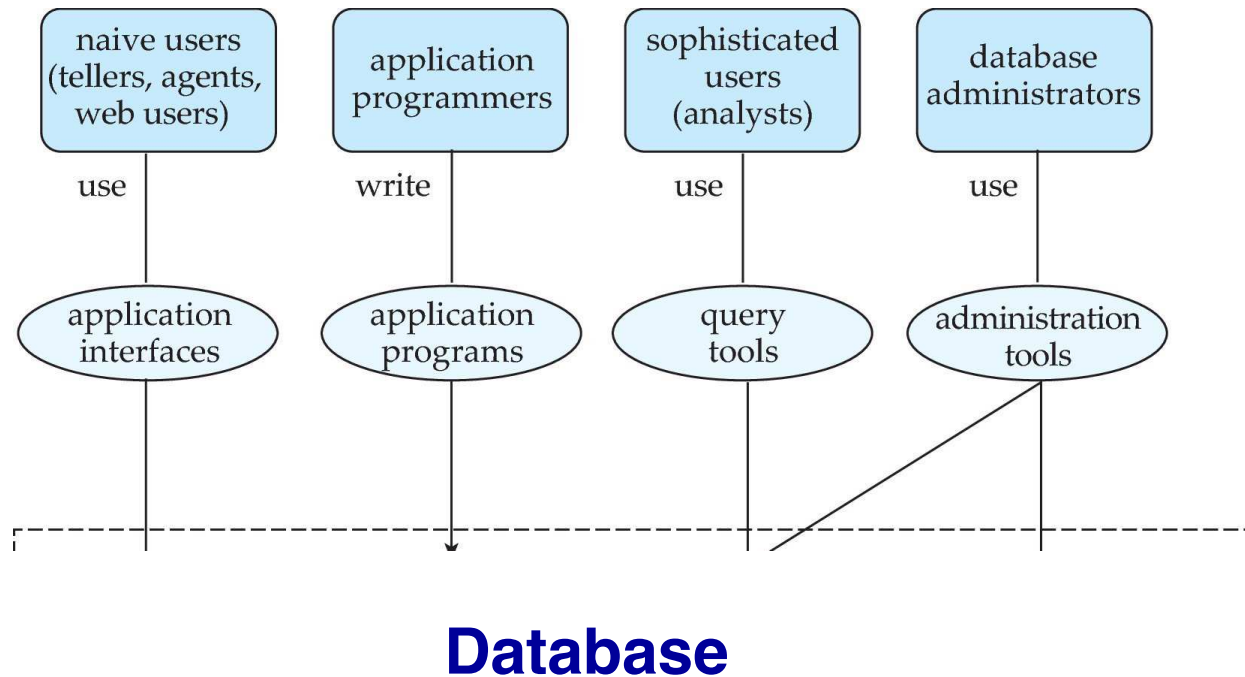
Database

Database System Internals

Database Architecture

The architecture of a database systems is greatly influenced by
 the underlying computer system on which the database is running:

- Centralized

- Client-server

- Parallel (multi-processor)

- Distributed

Build a Complete Database System in your free time?

- How much time do you need?

- To get a rough idea:
  - Postgres (about 800,000 lines of code)
    - Hundreds of man-years of work
  - Oracle (about 8,000,000 lines of code)
    - Probably thousands of man-years of work?

- Hmm, … probably not!

- Maybe a limited research prototype or new feature ;-)

History of Database Systems

- 1950s and early 1960s:
  - Data processing using magnetic tapes for storage
    - Tapes provided only sequential access
  - Punched cards for input

- Late 1960s and 1970s:
  - Hard disks allowed direct access to data
  - Network and hierarchical data models in widespread use
  - Ted Codd defines the relational data model
    - Would win the ACM Turing Award for this work
    - IBM Research begins System R prototype
    - UC Berkeley begins Ingres prototype
  - High-performance (for the era) transaction processing

History (cont.)

- 1980s:
  - Research relational prototypes evolve into commercial systems
    - SQL becomes industrial standard
  - Parallel and distributed database systems
  - Object-oriented database systems
- 1990s:
  - Large decision support and data-mining applications
  - Large multi-terabyte data warehouses
  - Emergence of Web commerce
- Early 2000s:
  - XML and XQuery standards
  - Automated database administration
- Later 2000s:
  - Giant data storage systems
    - Google BigTable, Yahoo PNuts, Amazon, ..

Recap

- Why databases?
- What do databases do?
- Data independence
  - Physical and Logical
- Database design
- Data models
  - Relational, object, XML, network, hierarchical
- Query languages
  - DML
  - DDL
- Architecture and systems aspects of database systems
  - Recovery
  - Concurrency control
  - Query processing (optimization)
  - File organization and indexing
- History of databases

End of Chapter 1

Outline

- Introduction
- **Relational Data Model**
- Formal Relational Languages (relational algebra)
- SQL
- Database Design
- Transaction Processing, Recovery, and Concurrency Control
- Storage and File Structures
- Indexing and Hashing
- Query Processing and Optimization

Figure 1.02

ID	name	dept_name	salary
22222	Einstein	Physics	95000
12121	Wu	Finance	90000
32343	El Said	History	60000
45565	Katz	Comp. Sci.	75000
98345	Kim	Elec. Eng.	80000
76766	Crick	Biology	72000
10101	Srinivasan	Comp. Sci.	65000
58583	Califieri	History	62000
83821	Brandt	Comp. Sci.	92000
15151	Mozart	Music	40000
33456	Gold	Physics	87000
76543	Singh	Finance	80000

(a) The *instructor* table

dept_name	building	budget
Comp. Sci.	Taylor	100000
Biology	Watson	90000
Elec. Eng.	Taylor	85000
Music	Packard	80000
Finance	Painter	120000
History	Painter	50000
Physics	Watson	70000

(b) The *department* table

Figure 1.04

| ID	name	salary	dept_name	building	budget
22222	Einstein	95000	Physics	Watson	70000
12121	Wu	90000	Finance	Painter	120000
32343	El Said	60000	History	Painter	50000
45565	Katz	75000	Comp. Sci.	Taylor	100000
98345	Kim	80000	Elec. Eng.	Taylor	85000
76766	Crick	72000	Biology	Watson	90000
10101	Srinivasan	65000	Comp. Sci.	Taylor	100000
58583	Califieri	62000	History	Painter	50000
83821	Brandt	92000	Comp. Sci	Taylor	100000
15151	Mozart	40000	Music	Packard	80000
33456	Gold	87000	Physics	Watson	70000
76543	Singh	80000	Finance	Painter	120000

Figure 1.06

(a) Two-tier architecture (b) Three-tier architecture

CS425 – Fall 2014
Boris Glavic
Chapter 2: Intro to Relational Model

Textbook: Chapter 2

Example of a Relation

attributes
(or columns)

ID	name	dept_name	salary
10101	Srinivasan	Comp. Sci.	65000
12121	Wu	Finance	90000
15151	Mozart	Music	40000
22222	Einstein	Physics	95000
32343	El Said	History	60000
33456	Gold	Physics	87000
45565	Katz	Comp. Sci.	75000
58583	Califieri	History	62000
76543	Singh	Finance	80000
76766	Crick	Biology	72000
83821	Brandt	Comp. Sci.	92000
98345	Kim	Elec. Eng.	80000

tuples
(or rows)

Attribute Types

- The set of allowed values for each attribute is called the **domain** or **data type** of the attribute

- Attribute values are (normally) required to be **atomic**; that is, indivisible
  - E.g., integer values
  - E.g., not address (street, city, zip code, state, country)

- The special value *null* is a member of every domain
  - Means *unknown* or *not applicable*

- The null value causes complications in the definition of many operations
  - Will be detailed later

Relation Schema and Instance

- A_1, A_2, …, A_n are **attributes names**

- $R = (A_1, A_2, …, A_n)$ is a **relation schema**

  Example:

  $instructor = (ID, name, dept\_name, salary)$

- Formally, given sets D_1, D_2, …. D_n of domains a **relation r** (or **relation instance)** is a subset of

  $D_1 \times D_2 \times … \times D_n$

  Thus, a relation is a **set** of **n-tuples** $(a_1, a_2, …, a_n)$ where each $a_i \in D_i$

- The current values (**relation instance**) of a relation are often specified in tabular form

  - Caveat: being a set, the tuples of the relation do not have any order defined as implied by the tabular representation

- An element **t** of **r** is a *tuple*, represented as a *row* in a table

Theory Warning

Alternative Definitions

- A relation schema is often defined as a list of attribute-domain pairs

  - That is the data types of each attribute in the relation are considered as part of the relation schema

- Tuples are sometimes defined as functions from attribute names to values (order of attributes does not matter)

- A relation **r** can be specified as a function

  - $D_1 \times D_2 \times \ldots \times D_n \to \{true, false\}$

  - $\mathbf{t} = (a_1, a_2, \ldots, a_n)$ is mapped to *true* if **t** is in **r** and to *false* otherwise

- These alternative definition are useful in database theory

  - We will stick to the simple definition!

Theory Warning

Relations are Unordered

- A relation is a **set** -> the elements of a set are not ordered per se
- From a pratical perspective:
  - Order of tuples is irrelevant (tuples may be stored in an arbitrary order)
- Example: *instructor* relation with unordered tuples

ID	name	dept_name	salary
22222	Einstein	Physics	95000
12121	Wu	Finance	90000
32343	El Said	History	60000
45565	Katz	Comp. Sci.	75000
98345	Kim	Elec. Eng.	80000
76766	Crick	Biology	72000
10101	Srinivasan	Comp. Sci.	65000
58583	Califieri	History	62000
83821	Brandt	Comp. Sci.	92000
15151	Mozart	Music	40000
33456	Gold	Physics	87000
76543	Singh	Finance	80000

Database

- A **database schema** *S* consists of multiple relation schema

- A **database instance** *I* for a schema S is a set of relation instances

  - One relation for each relation schema in S

- Information about an enterprise is broken up into parts

    instructor
    student
    advisor

- Bad design:
      univ (*instructor -ID, name, dept_name, salary, student_Id*, ..)
  results in

  - repetition of information (e.g., two students have the same instructor)

  - the need for many null values (e.g., represent an student with no advisor)

- Normalization theory (Chapter 7) deals with how to design "good" relational schemas avoiding these problems

Bad Design Example Revisited

- **Example:** Changing the budget of the 'Physics' department

  - Updates to many rows!

    - Easy to break **integrity**

    - If we forget to update a row, then we have multiple budget values for the physics department!

- Example: Deleting all employees from the 'Physics' department

  - How to avoid deleting the 'Physics' department?

  - Dummy employee's to store departments?

    - This is bad. E.g., counting the number of employees per department becomes more involved

ID	name	salary	dept_name	building	budget
22222	Einstein	95000	Physics	Watson	70000
12121	Wu	90000	Finance	Painter	120000
32343	El Said	60000	History	Painter	50000
45565	Katz	75000	Comp. Sci.	Taylor	100000
98345	Kim	80000	Elec. Eng.	Taylor	85000
76766	Crick	72000	Biology	Watson	90000
10101	Srinivasan	65000	Comp. Sci.	Taylor	100000
58583	Califieri	62000	History	Painter	50000
83821	Brandt	92000	Comp. Sci	Taylor	100000
15151	Mozart	40000	Music	Packard	80000
33456	Gold	87000	Physics	Watson	70000
76543	Singh	80000	Finance	Painter	120000

Keys

- Let $K \subseteq R$

- K is a **superkey** of R if values for K are sufficient to identify a unique tuple of each possible relation $r(R)$

  - Example: {ID} and {ID,name} are both superkeys of *instructor.*

- Superkey K is a **candidate key** if K is minimal (no subset of K is also a superkey)

  Example: {ID} is a candidate key for *Instructor*

- One of the candidate keys is selected to be the **primary key**.

  - which one? -> domain specific design choice

- **Foreign key** constraint: Value in one relation must appear in another

  - **Referencing** relation
 - **Referenced** relation

Keys

- Formally, a set of attributes $K \subseteq R$ is a superkey if for every instance r of R holds that

  - $\forall t, t' \in r: t.K = t'.K \Rightarrow t = t'$

- A superkey *K* is called a candidate key iff

  - $\forall K' \subseteq K: K'$ is not a superkey

- A foreign key constraint FK is quartuple (R, K, R', K') where R and R' are relation schemata, $K \subseteq R$, K' is the primary key of R', and |K| = |K'|

- A foreign key holds over an instance {r, r'} for {R,R'} iff

  - $\forall t \in R: \exists t' \in R': t.K = t'.K'$

Theory Warning

Schema Diagram for University Database

Relational Query Languages

- Procedural vs non-procedural (**declarative**)

- "Pure" languages:
  - Relational algebra
  - Tuple relational calculus
  - Domain relational calculus

- Expressive power of a query language
  - What queries can be expressed in this language?

- Relational algebra:
  - Algebra of relations -> set of operators that take relations as input and produce relations as output
  - -> **composable**: the output of evaluating an expression in relational algebra can be used as input to another relational algebra expression

- Now: First introduction to operators of the relational algebra

Selection of tuples

■ Relation r

A	B	C	D
α	α	1	7
α	β	5	7
β	β	12	3
β	β	23	10

■ Select tuples with A=B and D > 5

　　■ $\sigma_{A=B \text{ and } D > 5}$ (r)

A	B	C	D
α	α	1	7
β	β	23	10

Selection of Columns (Attributes)

■ Relation *r*:

A	B	C
α	10	1
α	20	1
β	30	1
β	40	2

■ Select A and C

  ■ Projection

  ■ $\Pi_{A,\,C}$ (r)

A	C
α	1
α	1
β	1
β	2

=

A	C
α	1
β	1
β	2

- Relations *r, s*:

A	B
α	1
β	2

r

C	D	E
α	10	a
β	10	a
β	20	b
γ	10	b

s

- *r* x *s*:

A	B	C	D	E
α	1	α	10	a
α	1	β	10	a
α	1	β	20	b
α	1	γ	10	b
β	2	α	10	a
β	2	β	10	a
β	2	β	20	b
β	2	γ	10	b

Union of two relations

- Relations *r, s:*

A	B
α	1
α	2
β	1

r

A	B
α	2
β	3

s

- r ∪ s:

A	B
α	1
α	2
β	1
β	3

Set difference of two relations

■ Relations r, s:

A	B
α	1
α	2
β	1

r

A	B
α	2
β	3

s

■ $r - s$:

A	B
α	1
β	1

Set Intersection of two relations

- Relation *r, s*:

A	B
α	1
α	2
β	1

r

A	B
α	2
β	3

s

- *r* ∩ *s*

A	B
α	2

Joining two relations – Natural Join

- Let r and s be relations on schemas R and S respectively. Then, the "natural join" of relations R and S is a relation on schema $R \cup S$ obtained as follows:

  - Consider each pair of tuples t_r from r and t_s from s.

  - If t_r and t_s have the same value on each of the attributes in $R \cap S$, add a tuple t to the result, where

    - t has the same value as t_r on r

    - t has the same value as t_s on s

Natural Join Example

- Relations r, s:

A	B	C	D
α	1	α	a
β	2	γ	a
γ	4	β	b
α	1	γ	a
δ	2	β	b

r

B	D	E
1	a	α
3	a	β
1	a	γ
2	b	δ
3	b	ε

s

- Natural Join
  - r ⋈ s

A	B	C	D	E
α	1	α	a	α
α	1	α	a	γ
α	1	γ	a	α
α	1	γ	a	γ
δ	2	β	b	δ

Figure in-2.1

Symbol (Name)	Example of Use
σ (Selection)	$\sigma_{\text{salary}>=85000}\ (instructor)$
	Return rows of the input relation that satisfy the predicate.
Π (Projection)	$\Pi_{ID,\ salary}\ (instructor)$
	Output specified attributes from all rows of the input relation. Remove duplicate tuples from the output.
⋈ (Natural Join)	$instructor \bowtie department$
	Output pairs of rows from the two input relations that have the same value on all attributes that have the same name.
× (Cartesian Product)	$instructor \times department$
	Output all pairs of rows from the two input relations (regardless of whether or not they have the same values on common attributes)
∪ (Union)	$\Pi_{name}(instructor) \cup \Pi_{name}(student)$
	Output the union of tuples from the two input relations.

End of Chapter 2

Recap

- **Database Schema** (or short schema)
  - Set of **relation schemata**
    - List of **attribute names**
- **Database Instance** (or short database)
  - Set of **relations instances**
    - Set of **tuples**
      - List of **attribute values**
- **Integrity Constraints**
  - **Keys** (Super-, Candidate-, Primary-)
    - For identifying tuples
  - **Foreign keys**
    - For referencing tuples in other relations
- **Query language**
  - **Declarative**
  - Retrieve, combine, and analyze data from a database instance

Outline

- Introduction
- Relational Data Model
- **Formal Relational Languages (relational algebra)**
- SQL
- Database Design
- Transaction Processing, Recovery, and Concurrency Control
- Storage and File Structures
- Indexing and Hashing
- Query Processing and Optimization

Figure 2.01

ID	name	dept_name	salary
10101	Srinivasan	Comp. Sci.	65000
12121	Wu	Finance	90000
15151	Mozart	Music	40000
22222	Einstein	Physics	95000
32343	El Said	History	60000
33456	Gold	Physics	87000
45565	Katz	Comp. Sci.	75000
58583	Califieri	History	62000
76543	Singh	Finance	80000
76766	Crick	Biology	72000
83821	Brandt	Comp. Sci.	92000
98345	Kim	Elec. Eng.	80000

Figure 2.02

course_id	title	dept_name	credits
BIO-101	Intro. to Biology	Biology	4
BIO-301	Genetics	Biology	4
BIO-399	Computational Biology	Biology	3
CS-101	Intro. to Computer Science	Comp. Sci.	4
CS-190	Game Design	Comp. Sci.	4
CS-315	Robotics	Comp. Sci.	3
CS-319	Image Processing	Comp. Sci.	3
CS-347	Database System Concepts	Comp. Sci.	3
EE-181	Intro. to Digital Systems	Elec. Eng.	3
FIN-201	Investment Banking	Finance	3
HIS-351	World History	History	3
MU-199	Music Video Production	Music	3
PHY-101	Physical Principles	Physics	4

Figure 2.03

course_id	prereq_id
BIO-301	BIO-101
BIO-399	BIO-101
CS-190	CS-101
CS-315	CS-101
CS-319	CS-101
CS-347	CS-101
EE-181	PHY-101

Figure 2.04

ID	name	dept_name	salary
22222	Einstein	Physics	95000
12121	Wu	Finance	90000
32343	El Said	History	60000
45565	Katz	Comp. Sci.	75000
98345	Kim	Elec. Eng.	80000
76766	Crick	Biology	72000
10101	Srinivasan	Comp. Sci.	65000
58583	Califieri	History	62000
83821	Brandt	Comp. Sci.	92000
15151	Mozart	Music	40000
33456	Gold	Physics	87000
76543	Singh	Finance	80000

Figure 2.05

dept_name	building	budget
Biology	Watson	90000
Comp. Sci.	Taylor	100000
Elec. Eng.	Taylor	85000
Finance	Painter	120000
History	Painter	50000
Music	Packard	80000
Physics	Watson	70000

Figure 2.06

course_id	sec_id	semester	year	building	room_number	time_slot_id
BIO-101	1	Summer	2009	Painter	514	B
BIO-301	1	Summer	2010	Painter	514	A
CS-101	1	Fall	2009	Packard	101	H
CS-101	1	Spring	2010	Packard	101	F
CS-190	1	Spring	2009	Taylor	3128	E
CS-190	2	Spring	2009	Taylor	3128	A
CS-315	1	Spring	2010	Watson	120	D
CS-319	1	Spring	2010	Watson	100	B
CS-319	2	Spring	2010	Taylor	3128	C
CS-347	1	Fall	2009	Taylor	3128	A
EE-181	1	Spring	2009	Taylor	3128	C
FIN-201	1	Spring	2010	Packard	101	B
HIS-351	1	Spring	2010	Painter	514	C
MU-199	1	Spring	2010	Packard	101	D
PHY-101	1	Fall	2009	Watson	100	A

Figure 2.07

ID	course_id	sec_id	semester	year
10101	CS-101	1	Fall	2009
10101	CS-315	1	Spring	2010
10101	CS-347	1	Fall	2009
12121	FIN-201	1	Spring	2010
15151	MU-199	1	Spring	2010
22222	PHY-101	1	Fall	2009
32343	HIS-351	1	Spring	2010
45565	CS-101	1	Spring	2010
45565	CS-319	1	Spring	2010
76766	BIO-101	1	Summer	2009
76766	BIO-301	1	Summer	2010
83821	CS-190	1	Spring	2009
83821	CS-190	2	Spring	2009
83821	CS-319	2	Spring	2010
98345	EE-181	1	Spring	2009

Figure 2.10

ID	name	dept_name	salary
12121	Wu	Finance	90000
22222	Einstein	Physics	95000
33456	Gold	Physics	87000
83821	Brandt	Comp. Sci.	92000

Figure 2.11

ID	salary
10101	65000
12121	90000
15151	40000
22222	95000
32343	60000
33456	87000
45565	75000
58583	62000
76543	80000
76766	72000
83821	92000
98345	80000

Figure 2.12

ID	name	salary	dept_name	building	budget
10101	Srinivasan	65000	Comp. Sci.	Taylor	100000
12121	Wu	90000	Finance	Painter	120000
15151	Mozart	40000	Music	Packard	80000
22222	Einstein	95000	Physics	Watson	70000
32343	El Said	60000	History	Painter	50000
33456	Gold	87000	Physics	Watson	70000
45565	Katz	75000	Comp. Sci.	Taylor	100000
58583	Califieri	62000	History	Painter	50000
76543	Singh	80000	Finance	Painter	120000
76766	Crick	72000	Biology	Watson	90000
83821	Brandt	92000	Comp. Sci.	Taylor	100000
98345	Kim	80000	Elec. Eng.	Taylor	85000

Figure 2.13

ID	salary
12121	90000
22222	95000
33456	87000
83821	92000

CS425 – Fall 2013
Boris Glavic
Chapter 3: Formal Relational Query Languages

Modified from:

Database System Concepts, 6th Ed.

Chapter 3: Formal Relational Query Languages

- Relational Algebra
- Tuple Relational Calculus
- Domain Relational Calculus

Textbook: Chapter 6

Relational Algebra

- Procedural language
- Six basic operators

  - select: σ
  - project: $\prod$
  - union: $\cup$
  - set difference: $-$
  - Cartesian product: x
  - rename: ρ

- The operators take one or two relations as inputs and produce a new relation as a result.

  - **composable**

Select Operation – Example

■ Relation r

$$
\begin{array}{|c|c|c|c|}
\hline
A & B & C & D \\
\hline
\alpha & \alpha & 1 & 7 \\
\alpha & \beta & 5 & 7 \\
\beta & \beta & 12 & 3 \\
\beta & \beta & 23 & 10 \\
\hline
\end{array}
$$

■ $\sigma_{A=B \wedge D > 5}(r)$

$$
\begin{array}{|c|c|c|c|}
\hline
A & B & C & D \\
\hline
\alpha & \alpha & 1 & 7 \\
\beta & \beta & 23 & 10 \\
\hline
\end{array}
$$

Select Operation

- Notation: $\sigma_p(r)$
- p is called the **selection predicate**
- Defined as:

$$\sigma_p(r) = \{t \mid t \in r \land p(t)\}$$

Where p is a formula in propositional calculus consisting of **terms** connected by : ∧ (**and**), ∨ (**or**), ¬ (**not**)
Each **term** is one of:

<attribute>   *op*   <attribute> or <constant>

where *op* is one of: $=, \neq, >, \geq. <. \leq$

- Example of selection:

$$\sigma_{dept\_name=\text{“Physics”}}(instructor)$$

Project Operation – Example

- Relation r:

A	B	C
α	10	1
α	20	1
β	30	1
β	40	2

- $\prod_{A,C} (r)$

A	C
α	1
α	1
β	1
β	2

$=$

A	C
α	1
β	1
β	2

Project Operation

- Notation:

$$\prod_{A_1, A_2, \ldots, A_k}(r)$$

  where A_1, A_2 are attribute names and r is a relation name.

- The result is defined as the relation of k columns obtained by erasing the columns that are not listed

- Duplicate rows removed from result, since relations are sets

- Let A be a subset of the attributes of relation r then:

$$\pi_A(r) = \{t.A \mid t \in r\}$$

- Example: To eliminate the *dept_name* attribute of *instructor*

$$\prod_{ID,\ name,\ salary} (instructor)$$

Theory Warning

Union Operation – Example

- Relations *r, s:*

A	B
α	1
α	2
β	1

r

A	B
α	2
β	3

s

- r ∪ s:

A	B
α	1
α	2
β	1
β	3

Union Operation

- Notation: $r \cup s$

- Defined as:

$$r \cup s = \{t \mid t \in r \lor t \in s\}$$

- For $r \cup s$ to be valid.

  1. r, s must have the *same* **arity** (same number of attributes)
  2. The attribute domains must be **compatible** (example: 2nd column of r deals with the same type of values as does the 2nd column of s)

- Example: to find all courses taught in the Fall 2009 semester, or in the Spring 2010 semester, or in both

$$\Pi_{course\_id}(\sigma_{semester=\text{"Fall"} \land year=2009}(section)) \cup$$
$$\Pi_{course\_id}(\sigma_{semester=\text{"Spring"} \land year=2010}(section))$$

Set difference of two relations

- Relations *r, s*:

A	B
α	1
α	2
β	1

r

A	B
α	2
β	3

s

- *r − s:*

A	B
α	1
β	1

Set Difference Operation

- Notation $r - s$

- Defined as:

$$r - s = \{t \mid t \in r \wedge t \notin s\}$$

- Set differences must be taken between **compatible** relations.
  - r and s must have the same arity
  - attribute domains of r and s must be compatible

- Example: to find all courses taught in the Fall 2009 semester, but not in the Spring 2010 semester

$$\Pi_{course\_id}\,(\sigma_{\,semester=\text{``Fall''}\,\wedge\,year=2009}\,(section)) \; -$$
$$\Pi_{course\_id}\,(\sigma_{\,semester=\text{``Spring''}\,\wedge\,year=2010}\,(section))$$

Cartesian-Product Operation – Example

- Relations *r, s*:

A	B
α	1
β	2

r

C	D	E
α	10	a
β	10	a
β	20	b
γ	10	b

s

- *r* x *s*:

A	B	C	D	E
α	1	α	10	a
α	1	β	10	a
α	1	β	20	b
α	1	γ	10	b
β	2	α	10	a
β	2	β	10	a
β	2	β	20	b
β	2	γ	10	b

Cartesian-Product Operation

- Notation *r* x *s*

- Defined as:

$$r \times s = \{t, t' \mid t \in r \wedge t' \in s\}$$

- Assume that attributes of r(R) and s(S) are disjoint. (That is, $R \cap S = \varnothing$).

- If attributes of *r(R)* and *s(S)* are not disjoint, then renaming must be used.

Composition of Operations

- Can build expressions using multiple operations

- Example: $\sigma_{A=C}(r \times s)$

- $r \times s$

A	B	C	D	E
α	1	α	10	a
α	1	β	10	a
α	1	β	20	b
α	1	γ	10	b
β	2	α	10	a
β	2	β	10	a
β	2	β	20	b
β	2	γ	10	b

- $\sigma_{A=C}(r \times s)$

A	B	C	D	E
α	1	α	10	a
β	2	β	10	a
β	2	β	20	b

Rename Operation

- Allows us to name, and therefore to refer to, the results of relational-algebra expressions.

- Allows us to refer to a relation by more than one name.

- Example:

$$\rho_X(r)$$

returns the expression E under the name X

- If a relational-algebra expression E has arity n, then

$$\rho_{x(A_1, A_2, \ldots, A_n)}(r)$$

returns the result of expression E under the name X, and with the

attributes renamed to $A_1, A_2, \ldots, A_n$.

$$\rho_X(r) = \{t(X) \mid t \in r\}$$

$$\rho_{X(A)}(r) = \{t(X).A \mid t \in r\}$$

Example Query

- Find the largest salary in the university
  - Step 1: find instructor salaries that are less than some other instructor salary (i.e. not maximum)
    - using a copy of *instructor* under a new name *d*

$$\pi_{instructor.salary}(\sigma_{instructor.salary < d.salary}$$
$$(instructor \times \rho_d(instructor)))$$

  - Step 2: Find the largest salary

$$\pi_{salary}(instructor)-$$
$$\pi_{instructor.salary}(\sigma_{instructor.salary < d.salary}$$
$$(instructor \times \rho_d(instructor)))$$

Example Queries

■ Find the names of all instructors in the Physics department, along with the *course_id* of all courses they have taught

- Query 1

$$\pi_{instructor.ID,course\_id}(\sigma_{dept\_name='Physics'}(\\ \sigma_{instructor.ID=teaches.ID}(instructor \times teaches)))$$

- Query 2

$$\pi_{instructor.ID,course\_id}(\sigma_{instructor.ID=teaches.ID}(\\ \sigma_{dept\_name='Physics'}(instructor \times teaches)))$$

Formal Definition (Syntax)

- A basic expression in the relational algebra consists of either one of the following:

  - A relation in the database

  - A constant relation: e.g., {(1),(2)}

- Let E_1 and E_2 be relational-algebra expressions; the following are all relational-algebra expressions:

  - $E_1 \cup E_2$

  - $E_1 - E_2$

  - $E_1 \times E_2$

  - $\sigma_p(E_1)$, P is a predicate on attributes in E_1

  - $\prod_s(E_1)$, S is a list consisting of some of the attributes in E_1

  - $\rho_x(E_1)$, x is the new name for the result of E_1

Theory Warning

Formal Definition (Semantics)

- Let E be an relational algebra expression. We use [E](I) to denote the evaluation of E over a database instance I

    - For simplicity we will often drop I and []

- The result of evaluating a simple relational algebra expression E over a database is defined as

    - Simple relation: [R](I) = R(I)

    - Constant relation: [C](I) = C

Theory Warning

Formal Definition (Semantics)

- Let E_1 and E_2 be relational-algebra expressions.

$$[E_1 \cup E_2] = \{t \mid t \in [E_1] \vee t \in [E_2]\}$$

$$[E_1 - E_2] = \{t \mid t \in [E_1] \wedge t \notin [E_2]\}$$

$$[E_1 \times E_2] = \{t, t' \mid t \in [E_1] \wedge t' \in [E_2]\}$$

$$[\sigma_p(E_1)] = \{t \mid t \in [E_1] \wedge p(t)\}$$

$$[\pi_A(E_1)] = \{t.A \mid t \in [E_1]\}$$

$$[\rho_X(E_1)] = \{t(X) \mid t \in [E_1]\}$$

Theory Warning

Null Values

- It is possible for tuples to have a null value, denoted by *null*, for some of their attributes

- *null* signifies an unknown value or that a value does not exist.

- The result of any arithmetic expression involving *null* is *null.*

- Aggregate functions simply ignore null values (as in SQL)

- For duplicate elimination and grouping, null is treated like any other value, and two nulls are assumed to be the same (as in SQL)

Null Values

- Comparisons with null values return the special truth value: *unknown*

  - If *false* was used instead of *unknown*, then   *not (A < 5)*
    would not be equivalent to   *A >= 5*

- Three-valued logic using the truth value *unknown*:

  - OR: (*unknown* **or** *true*)   = *true*,
    (*unknown* **or** *false*)   = *unknown*
    (*unknown* **or** *unknown*) = *unknown*

  - AND: (*true* **and** *unknown*)   = *unknown,*
    (*false* **and** *unknown*)   = *false,*
    (*unknown* **and** *unknown*) = *unknown*

  - NOT*:* (**not** *unknown*) = *unknown*

  - In SQL "*P* **is unknown**" evaluates to true if predicate *P* evaluates to *unknown*

- Result of select predicate is treated as *false* if it evaluates to *unknown*

Additional Operations

We define additional operations that do not add any expressive power to the relational algebra, but that simplify common queries.

- Set intersection

- Natural join

- Assignment

- Outer join

Set-Intersection Operation

- Notation: $r \cap s$

- Defined as:

$$r \cap s = \{t \mid t \in r \wedge t \in s\}$$

- Assume:
  - *r*, *s* have the *same arity*
  - attributes of *r* and *s* are compatible

- Note: $r \cap s = r - (r - s)$
  - That is adding intersection to the language does not make it more expressive

Set-Intersection Operation – Example

- Relation *r, s*:

A	B
α	1
α	2
β	1

r

A	B
α	2
β	3

s

- *r* ∩ *s*

A	B
α	2

Natural-Join Operation

- Notation: $r \bowtie s$

- Let r and s be relations on schemas R and S respectively.
  Then, $r \bowtie s$ is a relation on schema $R \cup S$ obtained as follows:

  - Consider each pair of tuples t_r from r and t_s from s.
  - If t_r and t_s have the same value on each of the attributes in $R \cap S$, add a tuple t to the result, where
    - t has the same value as t_r on r
    - t has the same value as t_s on s

- Example:

  $R = (A, B, C, D)$

  $S = (E, B, D)$

  - Result schema = (A, B, C, D, E)
  - $r \bowtie s$ is defined as:

  $$\prod_{r.A,\ r.B,\ r.C,\ r.D,\ s.E} (\sigma_{r.B = s.B \,\wedge\, r.D = s.D} (r \times s))$$

■ Let *r* and *s* be relations on schemas *R* and *S* respectively. Then, r ⋈ s is defined as:

$$X = R \cap S$$

$$S' = S - R$$

$$r \bowtie s = \pi_{R,S'}\big(\sigma_{r.X=s.X}(r \times s)\big)$$

Natural Join Example

■ Relations r, s:

A	B	C	D
α	1	α	a
β	2	γ	a
γ	4	β	b
α	1	γ	a
δ	2	β	b

r

B	D	E
1	a	α
3	a	β
1	a	γ
2	b	δ
3	b	ε

s

■ r ⋈ s

A	B	C	D	E
α	1	α	a	α
α	1	α	a	γ
α	1	γ	a	α
α	1	γ	a	γ
δ	2	β	b	δ

Natural Join and Theta Join

- Find the names of all instructors in the Comp. Sci. department together with the course titles of all the courses that the instructors teach

  - $\prod_{name,\ title} (\sigma_{dept\_name=\text{"Comp. Sci."}} (instructor \bowtie teaches \bowtie course))$

- Natural join is associative

  - $(instructor \bowtie teaches) \bowtie course$ is equivalent to $instructor \bowtie (teaches \bowtie course)$

- Natural join is commutative (we ignore attribute order)

  - $instruct \bowtie teaches$ is equivalent to $teaches \bowtie instructor$

- The **theta join** operation $r \bowtie_\theta s$ is defined as

$$r \bowtie_\theta s = \sigma_\theta(r \times s)$$

Assignment Operation

- The assignment operation ($\leftarrow$) provides a convenient way to express complex queries.

  - Write query as a sequential program consisting of
    - a series of assignments
    - followed by an expression whose value is displayed as a result of the query.
  - Assignment must always be made to a temporary relation variable.

$$E_1 \leftarrow \sigma_{salary > 40000}(instructor)$$

$$E_2 \leftarrow \sigma_{salary < 10000}(instructor)$$

$$E_3 \leftarrow E_1 \cup E_2$$

Outer Join

- An extension of the join operation that avoids loss of information.

- Computes the join and then adds tuples form one relation that does not match tuples in the other relation to the result of the join.

- Uses *null* values:

  - *null* signifies that the value is unknown or does not exist

  - All comparisons involving *null* are (roughly speaking) **false** by definition.

    - We shall study precise meaning of comparisons with nulls later

Outer Join – Example

■ Relation *instructor1*

ID	name	dept_name
10101	Srinivasan	Comp. Sci.
12121	Wu	Finance
15151	Mozart	Music

■ Relation *teaches1*

ID	course_id
10101	CS-101
12121	FIN-201
76766	BIO-101

Outer Join – Example

- Join

instructor ⋈ *teaches*

ID	name	dept_name	course_id
10101	Srinivasan	Comp. Sci.	CS-101
12121	Wu	Finance	FIN-201

- Left Outer Join

instructor ⟕ *teaches*

ID	name	dept_name	course_id
10101	Srinivasan	Comp. Sci.	CS-101
12121	Wu	Finance	FIN-201
15151	Mozart	Music	*null*

Outer Join – Example

- Right Outer Join

  instructor ⟕ *teaches*

ID	name	dept_name	course_id
10101	Srinivasan	Comp. Sci.	CS-101
12121	Wu	Finance	FIN-201
76766	null	null	BIO-101

- Full Outer Join

  instructor ⟗ *teaches*

ID	name	dept_name	course_id
10101	Srinivasan	Comp. Sci.	CS-101
12121	Wu	Finance	FIN-201
15151	Mozart	Music	*null*
76766	null	null	BIO-101

Outer Join using Joins

- Outer join can be expressed using basic operations

$$r \mathrel{\rlap{\sqsupset}{\bowtie}} s = (r \bowtie s) \cup ((r - \Pi_R(r \bowtie s)) \times \{(null, \dots, null)\})$$

$$r \mathrel{\bowtie\rlap{\sqsubset}{}} s = (r \bowtie s) \cup (\{(null, \dots, null)\} \times (s - \Pi_S(r \bowtie s)))$$

$$r \mathrel{\rlap{\sqsupset}{\bowtie}\rlap{\sqsubset}{}} s = (r \bowtie s) \cup ((r - \Pi_R(r \bowtie s)) \times \{(null, \dots, null)\})$$
$$\cup (\{(null, \dots, null)\} \times (s - \Pi_S(r \bowtie s)))$$

Division Operator

- Given relations r(R) and s(S), such that S $\subset$ R, r $\div$ s is the largest relation t(R-S) such that

  $$t \times s \subseteq r$$

  - Alternatively, all tuples from r.(R-S) such that all their extensions on R $\cap$ S with tuples from s exist in R

- E.g. let $r(ID, course\_id) = \prod_{ID,\ course\_id} (takes)$ and

  $s(course\_id) = \prod_{course\_id} (\sigma_{dept\_name=\text{"Biology"}}(course))$

  then r $\div$ s gives us students who have taken all courses in the Biology department

- Can write r $\div$ s as

$$E_1 \leftarrow \Pi_{R-S}(r)$$

$$E_2 \leftarrow \Pi_{R-S}((E_1 \times s) - \Pi_{R-S,S}(r \bowtie s))$$

$$r \div s = E_1 - E_2$$

Division Operator Example

■ Return the name of all persons that read all newspapers

reads

name	newspaper
Peter	Times
Bob	Wall Street
Alice	Times
Alice	Wall Street

newspaper

newspaper
Times
Wall Street

$$E_1 \leftarrow \Pi_{name}(reads)$$

$$E_2 \leftarrow ((E_1 \times newspaper) - \Pi_{name, newspaper}(reads \bowtie newspaper))$$

$$reads \div newspaper = E_1 - E_2$$

$$[reads \div newspaper] = \{(Alice)\}$$

Extended Relational-Algebra-Operations

- Generalized Projection
- Aggregate Functions

Generalized Projection

■ Extends the projection operation by allowing arithmetic functions to be used in the projection list.

$$\pi_{F_1,\ldots,F_n}(E)$$

■ *E* is any relational-algebra expression

■ Each of F_1, F_2, …, F_n are arithmetic expressions and function calls involving constants and attributes in the schema of *E*.

■ Given relation *instructor(ID, name, dept_name,* salary) where salary is annual salary, get the same information but with monthly salary

$$\prod_{ID,\ name,\ dept\_name,\ salary/12}(instructor)$$

■ *Adding functions increases expressive power!*

   ● *In standard relational algebra there is no way to change attribute values*

Aggregate Functions and Operations

- **Aggregation function** takes a set of values and returns a single value as a result.

  avg: average value
  min: minimum value
  max: maximum value
  sum: sum of values
  count: number of values

- **Aggregate operation** in relational algebra

$$G_1, G_2, \ldots, G_n \, \mathcal{G}_{F_1(A_1), F_2(A_2), \ldots, F_n(A_n)}(E)$$

  E is any relational-algebra expression

  - $G_1, G_2, \ldots, G_n$ is a list of attributes on which to group (can be empty)
  - Each F_i is an aggregate function
  - Each A_i is an attribute name

- Note: Some books/articles use γ instead of $\mathcal{G}$ (Calligraphic G)

Aggregate Operation – Example

- Relation r:

A	B	C
α	α	7
α	β	7
β	β	3
β	β	10

- $\mathcal{G}_{\textbf{sum(c)}}(r)$

sum(c)
27

Aggregate Operation – Example

- Find the average salary in each department

$$_{dept\_name}\,\mathcal{G}\,_{\textbf{avg}(salary)}\,(instructor)$$

ID	name	dept_name	salary
76766	Crick	Biology	72000
45565	Katz	Comp. Sci.	75000
10101	Srinivasan	Comp. Sci.	65000
83821	Brandt	Comp. Sci.	92000
98345	Kim	Elec. Eng.	80000
12121	Wu	Finance	90000
76543	Singh	Finance	80000
32343	El Said	History	60000
58583	Califieri	History	62000
15151	Mozart	Music	40000
33456	Gold	Physics	87000
22222	Einstein	Physics	95000

dept_name	avg_salary
Biology	72000
Comp. Sci.	77333
Elec. Eng.	80000
Finance	85000
History	61000
Music	40000
Physics	91000

Aggregate Functions (Cont.)

- What are the names for attributes in aggregation results?
  - Need some convention!
    - ▸ E.g., use the expression as a name **avg**(salary)
  - For convenience, we permit renaming as part of aggregate operation

$$_{dept\_name}\,\mathcal{G}\,_{\textbf{avg}(salary)\,\textbf{\textit{as}}\,avg\_sal}\,(instructor)$$

Modification of the Database

- The content of the database may be modified using the following operations:

  - Deletion

  - Insertion

  - Updating

- All these operations can be expressed using the assignment operator

- Example: Delete instructors with salary over $1,000,000

$$R \leftarrow R - (\sigma_{salary > 1000000}(R))$$

3.44

Restrictions for Modification

- Consider a modification where R=(A,B) and S=(C)

$$R \leftarrow \sigma_{C>5}(S)$$

- This would change the schema of R!
  - Should not be allowed

- Requirements for modifications
  - The name **R** on the left-hand side of the assignment operator refers to an existing relation in the database schema
  - The expression on the right-hand side of the assignment operator should be union-compatible with **R**

Tuple Relational Calculus

Tuple Relational Calculus

■ A nonprocedural query language, where each query is of the form

$$\{t \mid P(t)\}$$

■ It is the set of all tuples t such that predicate P is true for t

■ t is a *tuple variable*, $t[A]$ denotes the value of tuple t on attribute A

■ $t \in r$ denotes that tuple t is in relation r

■ P is a *formula* similar to that of the predicate calculus

Predicate Calculus Formula

1. Set of attributes and constants

2. Set of comparison operators: (e.g., $<$, $\leq$, $=$, $\neq$, $>$, $\geq$)

3. Set of logical connectives: and ($\wedge$), or ($\vee$), not ($\neg$)

4. Implication ($\Rightarrow$): x $\Rightarrow$ y, if x if true, then y is true

$$x \Rightarrow y \equiv \neg x \vee y$$

5. Set of quantifiers:

   ▸ $\exists\, t \in r\, (Q\,(t\,)) \equiv$ "there exists" a tuple in t in relation r
   such that predicate $Q\,(t\,)$ is true

   ▸ $\forall t \in r\, (Q\,(t\,)) \equiv Q$ is true "for all" tuples t in relation r

Example Queries

■ Find the *ID, name, dept_name, salary* for instructors whose salary is greater than $80,000

$$\{t \mid t \in instructor \wedge t\,[salary\,] > 80000\}$$

■ As in the previous query, but output only the *ID* attribute value

$$\{t \mid \exists\, s \in instructor\,(t\,[ID\,] = s\,[ID\,] \wedge s\,[salary\,] > 80000)\}$$

Notice that a relation on schema (*ID*) is implicitly defined by

the query, because

      1) t is not bound to any relation by the predicate

      2) we implicitly state that t has an ID attribute (*t[ID] = s[ID]*)

Example Queries

- Find the names of all instructors whose department is in the Watson building

$$\{t \mid \exists s \in instructor \, (t\,[name\,] = s\,[name\,] \\ \wedge \, \exists u \in department \, (u\,[dept\_name\,] = s[dept\_name]\, " \\ \wedge \, u\,[building] = \text{``Watson''} \,))\}$$

- Find the set of all courses taught in the Fall 2009 semester, or in the Spring 2010 semester, or both

$$\{t \mid \exists s \in section \, (t\,[course\_id\,] = s\,[course\_id\,] \wedge \\ s\,[semester] = \text{``Fall''} \wedge s\,[year] = 2009) \\ \vee \, \exists u \in section \, (t\,[course\_id\,] = u\,[course\_id\,] \wedge \\ u\,[semester] = \text{``Spring''} \wedge u\,[year] = 2010)\}$$

Example Queries

- Find the set of all courses taught in the Fall 2009 semester, and in the Spring 2010 semester

    $\{t \mid \exists s \in section \, (t\,[course\_id\,] = s\,[course\_id\,] \,\wedge$
    $s\,[semester] = \text{``Fall''} \wedge s\,[year] = 2009\,)$
    $\wedge \; \exists u \in section \, (t\,[course\_id\,] = u\,[course\_id\,] \,\wedge$
    $u\,[semester] = \text{``Spring''} \wedge u\,[year] = 2010)\}$

- Find the set of all courses taught in the Fall 2009 semester, but not in the Spring 2010 semester

    $\{t \mid \exists s \in section \, (t\,[course\_id\,] = s\,[course\_id\,] \,\wedge$
    $s\,[semester] = \text{``Fall''} \wedge s\,[year] = 2009)$
    $\wedge \; \neg \; \exists u \in section \, (t\,[course\_id\,] = u\,[course\_id\,] \,\wedge$
    $u\,[semester] = \text{``Spring''} \wedge u\,[year] = 2010)\}$

Safety of Expressions

■ It is possible to write tuple calculus expressions that generate infinite relations.

■ For example, { t | ¬ t ∈ r } results in an infinite relation if the domain of any attribute of relation r is infinite

■ To guard against the problem, we restrict the set of allowable expressions to safe expressions.

■ An expression {t | P (t)} in the tuple relational calculus is *safe* if every component of t appears in one of the relations, tuples, or constants that appear in P

  ● NOTE: this is more than just a syntax condition.

    ‣ E.g. { t | t [A] = 5 ∨ **true** } is not safe --- it defines an infinite set with attribute values that do not appear in any relation or tuples or constants in P.

Universal Quantification

■ Find all students who have taken all courses offered in the Biology department

● $\{t \mid \exists\ r \in student\ (t\ [ID] = r\ [ID])\ \wedge$
$(\forall\ u \in course\ (u\ [dept\_name]="Biology" \Rightarrow$
$\exists\ s \in takes\ (t\ [ID] = s\ [ID\ ] \wedge$
$s\ [course\_id] = u\ [course\_id]))\}$

● Note that without the existential quantification on student, the above query would be unsafe if the Biology department has not offered any courses.

Domain Relational Calculus

Domain Relational Calculus

■ A nonprocedural query language equivalent in power to the tuple relational calculus

■ Each query is an expression of the form:

$$\{ <x_1, x_2, \ldots, x_n> \mid P(x_1, x_2, \ldots, x_n) \}$$

- $x_1, x_2, \ldots, x_n$ represent domain variables
  - ▸ Variables that range of attribute values
- P represents a formula similar to that of the predicate calculus
- Tuples can be formed using <>
  - ▸ E.g., <'Einstein','Physics'>

Example Queries

- Find the *ID, name, dept_name, salary* for instructors whose salary is greater than $80,000

  - $\{<i, n, d, s> |\ <i, n, d, s> \in instructor \wedge s > 80000\}$

- As in the previous query, but output only the *ID* attribute value

  - $\{<i> |\ <i, n, d, s> \in instructor \wedge s > 80000\}$

- Find the names of all instructors whose department is in the Watson building

  $\{<n> |\ \exists\ i, d, s\ (<i, n, d, s> \in instructor$
  $\wedge\ \exists\ b, a\ (<d, b, a> \in department\ \wedge\ b = \text{``Watson''}\ ))\}$

Example Queries

■ Find the set of all courses taught in the Fall 2009 semester, or in the Spring 2010 semester, or both

$$\{<c> | \, \exists \, a, s, y, b, r, t \, (\, <c, a, s, y, b, t> \in section \, \wedge$$
$$s = \text{``Fall''} \wedge y = 2009 \,)$$
$$\vee \, \exists \, a, s, y, b, r, t \, (<c, a, s, y, b, t> \in section \,] \, \wedge$$
$$s = \text{``Spring''} \wedge y = 2010) \}$$

This case can also be written as
$$\{<c> | \, \exists \, a, s, y, b, r, t \, (\, <c, a, s, y, b, t> \in section \, \wedge$$
$$(\, (s = \text{``Fall''} \wedge y = 2009 \, ) \, \vee (s = \text{``Spring''} \wedge y = 2010)) \}$$

■ Find the set of all courses taught in the Fall 2009 semester, and in the Spring 2010 semester

$$\{<c> | \, \exists \, a, s, y, b, r, t \, (\, <c, a, s, y, b, t> \in section \, \wedge$$
$$s = \text{``Fall''} \wedge y = 2009 \,)$$
$$\wedge \, \exists \, a, s, y, b, r, t \, (<c, a, s, y, b, t> \in section \,] \, \wedge$$
$$s = \text{``Spring''} \wedge y = 2010) \}$$

Safety of Expressions

The expression:

$$\{ <x_1, x_2, \ldots, x_n> \mid P(x_1, x_2, \ldots, x_n) \}$$

is safe if all of the following hold:

1. All values that appear in tuples of the expression are values from *dom* (*P*) (that is, the values appear either as constants in *P* or in a tuple of a relation mentioned in *P*).

2. For every "there exists" subformula of the form $\exists x (P_1(x))$, the subformula is true if and only if there is a value of x in *dom* (P_1) such that $P_1(x)$ is true.

3. For every "for all" subformula of the form $\forall_x (P_1(x))$, the subformula is true if and only if $P_1(x)$ is true for all values x from *dom* (P_1).

Universal Quantification

■ Find all students who have taken all courses offered in the Biology department

- $\{< i > \mid \exists\ n, d, tc\ (< i, n, d, tc > \in student\ \wedge$
  $(\forall\ ci, ti, dn, cr\ (< ci, ti, dn, cr > \in course\ \wedge\ dn = \text{"Biology"}$
  $\Rightarrow\ \exists\ si, se, y, g\ (<i, ci, si, se, y, g> \in takes\ ))\}$

- Note that without the existential quantification on student, the above query would be unsafe if the Biology department has not offered any courses.

* Above query fixes bug in page 246, last query

Relationship between Relational Algebra and Tuple (Domain) Calculus

- **Codd's theorem**

  - Relational algebra and tuple calculus are equivalent

- That means that every query expressible in relational algebra can also be expressed in tuple calculus and vice versa

- Since domain calculus is as expressive as tuple calculus the same holds for the domain calculus

- Note: Here relational algebra refers to the standard version (no aggregation and projection with functions)

Recap

- **Relational algebra**
  - Standard relational algebra:
    - ▸ Selection, projection, renaming, cross product, union, set difference
  - Null values
  - Semantic sugar operators:
    - ▸ Intersection, joins, division,
  - Extensions:
    - ▸ Aggregation, extended projection
- **Tuple Calculus**
  - safety
- **Domain Calculus**

End of Chapter 3

Outline

- Introduction
- Relational Data Model
- Formal Relational Languages (relational algebra)
- **SQL - Introduction**
- Database Design
- Transaction Processing, Recovery, and Concurrency Control
- Storage and File Structures
- Indexing and Hashing
- Query Processing and Optimization

Figure 6.01

ID	name	dept_name	salary
10101	Srinivasan	Comp. Sci.	65000
12121	Wu	Finance	90000
15151	Mozart	Music	40000
22222	Einstein	Physics	95000
32343	El Said	History	60000
33456	Gold	Physics	87000
45565	Katz	Comp. Sci.	75000
58583	Califieri	History	62000
76543	Singh	Finance	80000
76766	Crick	Biology	72000
83821	Brandt	Comp. Sci.	92000
98345	Kim	Elec. Eng.	80000

Figure 6.02

ID	name	dept_name	salary
22222	Einstein	Physics	95000
33456	Gold	Physics	87000

Figure 6.03

ID	name	salary
10101	Srinivasan	65000
12121	Wu	90000
15151	Mozart	40000
22222	Einstein	95000
32343	El Said	60000
33456	Gold	87000
45565	Katz	75000
58583	Califieri	62000
76543	Singh	80000
76766	Crick	72000
83821	Brandt	92000
98345	Kim	80000

Figure 6.04

course_id	sec_id	semester	year	building	room_number	time_slot_id
BIO-101	1	Summer	2009	Painter	514	B
BIO-301	1	Summer	2010	Painter	514	A
CS-101	1	Fall	2009	Packard	101	H
CS-101	1	Spring	2010	Packard	101	F
CS-190	1	Spring	2009	Taylor	3128	E
CS-190	2	Spring	2009	Taylor	3128	A
CS-315	1	Spring	2010	Watson	120	D
CS-319	1	Spring	2010	Watson	100	B
CS-319	2	Spring	2010	Taylor	3128	C
CS-347	1	Fall	2009	Taylor	3128	A
EE-181	1	Spring	2009	Taylor	3128	C
FIN-201	1	Spring	2010	Packard	101	B
HIS-351	1	Spring	2010	Painter	514	C
MU-199	1	Spring	2010	Packard	101	D
PHY-101	1	Fall	2009	Watson	100	A

Figure 6.05

course_id
CS-101
CS-315
CS-319
CS-347
FIN-201
HIS-351
MU-199
PHY-101

Figure 6.06

course_id
CS-347
PHY-101

Figure 6.07

ID	course_id	sec_id	semester	year
10101	CS-101	1	Fall	2009
10101	CS-315	1	Spring	2010
10101	CS-347	1	Fall	2009
12121	FIN-201	1	Spring	2010
15151	MU-199	1	Spring	2010
22222	PHY-101	1	Fall	2009
32343	HIS-351	1	Spring	2010
45565	CS-101	1	Spring	2010
45565	CS-319	1	Spring	2010
76766	BIO-101	1	Summer	2009
76766	BIO-301	1	Summer	2010
83821	CS-190	1	Spring	2009
83821	CS-190	2	Spring	2009
83821	CS-319	2	Spring	2010
98345	EE-181	1	Spring	2009

Figure 6.08

Inst.ID	name	dept_name	salary	teaches.ID	course_id	sec_id	semester	year
10101	Srinivasan	Comp. Sci.	65000	10101	CS-101	1	Fall	2009
10101	Srinivasan	Comp. Sci.	65000	10101	CS-315	1	Spring	2010
10101	Srinivasan	Comp. Sci.	65000	10101	CS-347	1	Fall	2009
10101	Srinivasan	Comp. Sci.	65000	12121	FIN-201	1	Spring	2010
10101	Srinivasan	Comp. Sci.	65000	15151	MU-199	1	Spring	2010
10101	Srinivasan	Comp. Sci.	65000	22222	PHY-101	1	Fall	2009
...	...	...	...	...	...	...	...	...
...	...	...	...	...	...	...	...	...
12121	Wu	Finance	90000	10101	CS-101	1	Fall	2009
12121	Wu	Finance	90000	10101	CS-315	1	Spring	2010
12121	Wu	Pinance	90000	10101	CS-347	1	Fall	2009
12121	Wu	Pinance	90000	12121	FIN-201	1	Spring	2010
12121	Wu	Finance	90000	15151	MU-199	1	Spring	2010
12121	Wu	Pinance	90000	22222	PHY-101	1	Fall	2009
...	...	...	...	...	...	...	...	...
...	...	...	...	...	...	...	...	...
15151	Mozart	Music	40000	10101	CS-101	1	Fall	2009
15151	Mozart	Music	40000	10101	CS-315	1	Spring	2010
15151	Mozart	Music	40000	10101	CS-347	1	Fall	2009
15151	Mozart	Music	40000	12121	FIN-201	1	Spring	2010
15151	Mozart	Music	40000	15151	MU-199	1	Spring	2010
15151	Mozart	Music	40000	22222	PHY-101	1	Fall	2009
...	...	...	...	...	...	...	...	...
...	...	...	...	...	...	...	...	...
22222	Einstein	Physics	95000	10101	CS-101	1	Fall	2009
22222	Einstein	Physics	95000	10101	CS-315	1	Spring	2010
22222	Einstein	Physics	95000	10101	CS-347	1	Fall	2009
22222	Einstein	Physics	95000	12121	FIN-201	1	Spring	2010
22222	Einstein	Physics	95000	15151	MU-199	1	Spring	2010
22222	Einstein	Physics	95000	22222	PHY-101	1	Fall	2009
...	...	...	...	...	...	...	...	...
...	...	...	...	...	...	...	...	...

Figure 6.09

inst.ID	name	dept_name	salary	teaches.ID	course_id	sec_id	semester	year
22222	Einstein	Physics	95000	10101	CS-437	1	Fall	2009
22222	Einstein	Physics	95000	10101	CS-315	1	Spring	2010
22222	Einstein	Physics	95000	12121	FIN-201	1	Spring	2010
22222	Einstein	Physics	95000	15151	MU-199	1	Spring	2010
22222	Einstein	Physics	95000	22222	PHY-101	1	Fall	2009
22222	Einstein	Physics	95000	32343	HIS-351	1	Spring	2010
…	…	…	…	…	…	…	…	…
…	…	…	…	…	…	…	…	…
33456	Gold	Physics	87000	10101	CS-437	1	Fall	2009
33456	Gold	Physics	87000	10101	CS-315	1	Spring	2010
33456	Gold	Physics	87000	12121	FIN-201	1	Spring	2010
33456	Gold	Physics	87000	15151	MU-199	1	Spring	2010
33456	Gold	Physics	87000	22222	PHY-101	1	Fall	2009
33456	Gold	Physics	87000	32343	HIS-351	1	Spring	2010
…	…	…	…	…	…	…	…	…
…	…	…	…	…	…	…	…	…

Figure 6.10

name	course_id
Einstein	PHY-101

Figure 6.11

salary
65000
90000
40000
60000
87000
75000
62000
72000
80000
92000

Figure 6.12

salary
95000

Figure 6.13

course_id
CS-101

Figure 6.14

ID	name	dept_name	salary	course_id	sec_id	semester	year
10101	Srinivasan	Comp. Sci.	65000	CS-101	1	Fall	2009
10101	Srinivasan	Comp. Sci.	65000	CS-315	1	Spring	2010
10101	Srinivasan	Comp. Sci.	65000	CS-347	1	Fall	2009
12121	Wu	Finance	90000	FIN-201	1	Spring	2010
15151	Mozart	Music	40000	MU-199	1	Spring	2010
22222	Einstein	Physics	95000	PHY-101	1	Fall	2009
32343	El Said	History	60000	HIS-351	1	Spring	2010
45565	Katz	Comp. Sci.	75000	CS-101	1	Spring	2010
45565	Katz	Comp. Sci.	75000	CS-319	1	Spring	2010
76766	Crick	Biology	72000	BIO-101	1	Summer	2009
76766	Crick	Biology	72000	BIO-301	1	Summer	2010
83821	Brandt	Comp. Sci.	92000	CS-190	1	Spring	2009
83821	Brandt	Comp. Sci.	92000	CS-190	2	Spring	2009
83821	Brandt	Comp. Sci.	92000	CS-319	2	Spring	2010
98345	Kim	Elec. Eng.	80000	EE-181	1	Spring	2009

Figure 6.15

name	course_id
Srinivasan	CS-101
Srinivasan	CS-315
Srinivasan	CS-347
Wu	FIN-201
Mozart	MU-199
Einstein	PHY-101
El Said	HIS-351
Katz	CS-101
Katz	CS-319
Crick	BIO-101
Crick	BIO-301
Brandt	CS-190
Brandt	CS-319
Kim	EE-181

Figure 6.16

name	title
Brandt	Game Design
Brandt	Image Processing
Katz	Image Processing
Katz	Intro. to Computer Science
Srinivasan	Intro. to Computer Science
Srinivasan	Robotics
Srinivasan	Database System Concepts

Figure 6.17

ID	name	dept_name	salary	course_id	sec_id	semester	year
10101	Srinivasan	Comp. Sci.	65000	CS-101	1	Fall	2009
10101	Srinivasan	Comp. Sci.	65000	CS-315	1	Spring	2010
10101	Srinivasan	Comp. Sci.	65000	CS-347	1	Fall	2009
12121	Wu	Finance	90000	FIN-201	1	Spring	2010
15151	Mozart	Music	40000	MU-199	1	Spring	2010
22222	Einstein	Physics	95000	PHY-101	1	Fall	2009
32343	El Said	History	60000	HIS-351	1	Spring	2010
33456	Gold	Physics	87000	null	null	null	null
45565	Katz	Comp. Sci.	75000	CS-101	1	Spring	2010
45565	Katz	Comp. Sci.	75000	CS-319	1	Spring	2010
58583	Califieri	History	62000	null	null	null	null
76543	Singh	Finance	80000	null	null	null	null
76766	Crick	Biology	72000	BIO-101	1	Summer	2009
76766	Crick	Biology	72000	BIO-301	1	Summer	2010
83821	Brandt	Comp. Sci.	92000	CS-190	1	Spring	2009
83821	Brandt	Comp. Sci.	92000	CS-190	2	Spring	2009
83821	Brandt	Comp. Sci.	92000	CS-319	2	Spring	2010
98345	Kim	Elec. Eng.	80000	EE-181	1	Spring	2009

Figure 6.18

ID	course_id	sec_id	semester	year	name	dept_name	salary
10101	CS-101	1	Fall	2009	Srinivasan	Comp. Sci.	65000
10101	CS-315	1	Spring	2010	Srinivasan	Comp. Sci.	65000
10101	CS-347	1	Fall	2009	Srinivasan	Comp. Sci.	65000
12121	FIN-201	1	Spring	2010	Wu	Finance	90000
15151	MU-199	1	Spring	2010	Mozart	Music	40000
22222	PHY-101	1	Fall	2009	Einstein	Physics	95000
32343	HIS-351	1	Spring	2010	El Said	History	60000
33456	*null*	*null*	*null*	*null*	Gold	Physics	87000
45565	CS-101	1	Spring	2010	Katz	Comp. Sci.	75000
45565	CS-319	1	Spring	2010	Katz	Comp. Sci.	75000
58583	*null*	*null*	*null*	*null*	Califieri	History	62000
76543	*null*	*null*	*null*	*null*	Singh	Finance	80000
76766	BIO-101	1	Summer	2009	Crick	Biology	72000
76766	BIO-301	1	Summer	2010	Crick	Biology	72000
83821	CS-190	1	Spring	2009	Brandt	Comp. Sci.	92000
83821	CS-190	2	Spring	2009	Brandt	Comp. Sci.	92000
83821	CS-319	2	Spring	2010	Brandt	Comp. Sci.	92000
98345	EE-181	1	Spring	2009	Kim	Elec. Eng.	80000

Figure 6.19

ID	name	dept_name	salary
76766	Crick	Biology	72000
45565	Katz	Comp. Sci.	75000
10101	Srinivasan	Comp. Sci.	65000
83821	Brandt	Comp. Sci.	92000
98345	Kim	Elec. Eng.	80000
12121	Wu	Finance	90000
76543	Singh	Finance	80000
32343	El Said	History	60000
58583	Califieri	History	62000
15151	Mozart	Music	40000
33456	Gold	Physics	87000
22222	Einstein	Physics	95000

Figure 6.20

dept_name	salary
Biology	72000
Comp. Sci.	77333
Elec. Eng.	80000
Finance	85000
History	61000
Music	40000
Physics	91000

Figure 6.21

name
Einstein
Crick
Gold

Deletion

- A delete request is expressed similarly to a query, except instead of displaying tuples to the user, the selected tuples are removed from the database.

- Can delete only whole tuples; cannot delete values on only particular attributes

- A deletion is expressed in relational algebra by:

$$r \leftarrow r - E$$

where r is a relation and E is a relational algebra query.

Deletion Examples

- Delete all account records in the Perryridge branch.

$$account \leftarrow account - \sigma_{branch\_name = \text{"Perryridge"}}(account)$$

- Delete all loan records with amount in the range of 0 to 50

$$loan \leftarrow loan - \sigma_{amount \geq 0 \text{ and } amount \leq 50}(loan)$$

- Delete all accounts at branches located in Needham.

$$r_1 \leftarrow \sigma_{branch\_city = \text{"Needham"}}(account \bowtie branch)$$

$$r_2 \leftarrow \prod_{account\_number, branch\_name, balance}(r_1)$$

$$r_3 \leftarrow \prod_{customer\_name, account\_number}(r_2 \bowtie depositor)$$

$$account \leftarrow account - r_2$$

$$depositor \leftarrow depositor - r_3$$

Insertion

- To insert data into a relation, we either:
  - specify a tuple to be inserted
  - write a query whose result is a set of tuples to be inserted
- in relational algebra, an insertion is expressed by:

$$r \leftarrow r \cup E$$

  where r is a relation and E is a relational algebra expression.
- The insertion of a single tuple is expressed by letting E be a constant relation containing one tuple.

Insertion Examples

■ Insert information in the database specifying that Smith has $1200 in account A-973 at the Perryridge branch.

$$account \leftarrow account \cup \{(\text{``A-973''}, \text{``Perryridge''}, 1200)\}$$

$$depositor \leftarrow depositor \cup \{(\text{``Smith''}, \text{``A-973''})\}$$

■ Provide as a gift for all loan customers in the Perryridge branch, a $200 savings account. Let the loan number serve as the account number for the new savings account.

$$r_1 \leftarrow (\sigma_{branch\_name = \text{``Perryridge''}}(borrower \bowtie loan))$$

$$account \leftarrow account \cup \prod_{loan\_number, branch\_name, 200} (r_1)$$

$$depositor \leftarrow depositor \cup \prod_{customer\_name, loan\_number} (r_1)$$

Updating

- A mechanism to change a value in a tuple without charging *all* values in the tuple

- Use the generalized projection operator to do this task

$$r \leftarrow \prod_{F_1, F_2, \ldots, F_l,} (r)$$

- Each F_i is either

  - the I^{th} attribute of r, if the I^{th} attribute is not updated, or,

  - if the attribute is to be updated F_i is an expression, involving only constants and the attributes of r, which gives the new value for the attribute

Update Examples

- Make interest payments by increasing all balances by 5 percent.

$$account \leftarrow \prod_{account\_number,\ branch\_name,\ balance\ *\ 1.05} (account)$$

- Pay all accounts with balances over $10,000 6 percent interest and pay all others 5 percent

$$account \leftarrow \prod_{account\_number,\ branch\_name,\ balance\ *\ 1.06} (\sigma_{BAL > 10000} (account))$$
$$\cup \prod_{account\_number,\ branch\_name,\ balance\ *\ 1.05} (\sigma_{BAL \leq 10000} (account))$$

Example Queries

- Find the names of all customers who have a loan and an account at bank.

$$\prod_{customer\_name} (borrower) \cap \prod_{customer\_name} (depositor)$$

- Find the name of all customers who have a loan at the bank and the loan amount

$$\prod_{customer\_name,\ loan\_number,\ amount} (borrower \bowtie loan)$$

3.91

Example Queries

- Find all customers who have an account from at least the "Downtown" and the Uptown" branches.

  - Query 1

    $$\prod_{customer\_name} (\sigma_{branch\_name = \text{"Downtown"}} (depositor \bowtie account)) \cap$$

    $$\prod_{customer\_name} (\sigma_{branch\_name = \text{"Uptown"}} (depositor \bowtie account))$$

  - Query 2

    $$\prod_{customer\_name,\ branch\_name} (depositor \bowtie account)$$
    $$\div \rho_{temp(branch\_name)} (\{(\text{"Downtown"}), (\text{"Uptown"})\})$$

  Note that Query 2 uses a constant relation.

Bank Example Queries

- Find all customers who have an account at all branches located in Brooklyn city.

$$\prod_{customer\_name,\ branch\_name} (depositor \bowtie account)$$
$$\div \prod_{branch\_name} (\sigma_{branch\_city\ =\ \text{"Brooklyn"}} (branch))$$

CS425 – Fall 2013
Boris Glavic
Chapter 4: Introduction to SQL

Modified from:

Database System Concepts, 6th Ed.

Chapter 4: Introduction to SQL

- Overview of the SQL Query Language

- Data Definition

- Basic Query Structure

- Additional Basic Operations

- Set Operations

- Null Values

- Aggregate Functions

- Nested Subqueries

- Modification of the Database

Textbook: Chapter 3

History

- IBM Sequel language developed as part of System R project at the IBM San Jose Research Laboratory

- Renamed Structured Query Language (SQL)

- ANSI and ISO standard SQL:
  - SQL-86, SQL-89, SQL-92
  - SQL:1999, SQL:2003, SQL:2008

- Commercial systems offer most, if not all, SQL-92 features, plus varying feature sets from later standards and special proprietary features.
  - Not all examples here may work one-to-one on your particular system.

Data Definition Language

The SQL **data-definition language (DDL)** allows the specification of information about relations, including:

- The schema for each relation.

- The domain of values associated with each attribute.

- Integrity constraints

- And as we will see later, also other information such as

  - The set of indices to be maintained for each relations.

  - Security and authorization information for each relation.

  - The physical storage structure of each relation on disk.

Domain Types in SQL

- **char(n).** Fixed length character string, with user-specified length *n*.
- **varchar(n).** Variable length character strings, with user-specified maximum length *n*.
- **int.** Integer (a finite subset of the integers that is machine-dependent).
- **smallint.** Small integer (a machine-dependent subset of the integer domain type).
- **numeric(p,d).** Fixed point number, with user-specified precision of *p* digits, with *n* digits to the right of decimal point.
- **real, double precision.** Floating point and double-precision floating point numbers, with machine-dependent precision.
- **float(n).** Floating point number, with user-specified precision of at least *n* digits.
- More are covered in Chapter 4.

Create Table Construct

- An SQL relation is defined using the **create table** command:

  create table r (A_1 D_1, A_2 D_2, ..., A_n D_n,

          (integrity-constraint$_1$),

          ...,

          (integrity-constraint$_k$))

  - *r* is the name of the relation
  - each *A_i* is an attribute name in the schema of relation *r*
  - *D_i* is the data type of values in the domain of attribute *A_i*

- Example:

  create table *instructor* (
      ID   **char**(5),
      name   **varchar**(20) **not null,**
      dept_name   **varchar**(20),
      salary   **numeric**(8,2))

- **insert into** *instructor* **values** ('10211', 'Smith', 'Biology', 66000);
- **insert into** *instructor* **values** ('10211', null, 'Biology', 66000);

Integrity Constraints in Create Table

- **not null**

- **primary key** ($A_1, ..., A_n$)

- **foreign key** ($A_m, ..., A_n$) **references** *r*

Example: Declare *ID* as the primary key for *instructor*

.

> **create table** *instructor* (
>   *ID*   **char**(5),
>   *name*   **varchar**(20) **not null,**
>   *dept_name*   **varchar**(20),
>   *salary*   **numeric**(8,2),
>   **primary key** (*ID*),
>   **foreign key** *(dept_name)* **references** *department)*

primary key declaration on an attribute automatically ensures **not null**

And a Few More Relation Definitions

- **create table** *student* (
   ID   **varchar**(5),
   name   **varchar**(20) not null,
   dept_name **varchar**(20),
   tot_cred   **numeric**(3,0),
   primary key (*ID*),
   foreign key *(dept_name)* **references** *department)*);

- **create table** *takes* (
   ID   **varchar**(5),
   course_id **varchar**(8),
   sec_id   **varchar**(8),
   semester **varchar**(6),
   year   **numeric**(4,0),
   grade   **varchar**(2),
   primary key *(ID, course_id, sec_id, semester, year)*,
   foreign key (*ID*) **references** *student,*
   foreign key (*course_id, sec_id, semester, year*) **references** *section*);

  - Note: *sec_id* can be dropped from primary key above, to ensure a student cannot be registered for two sections of the same course in the same semester

Even more

- **create table** *course* (

      course_id   **varchar**(8) **primary key**,
      title   **varchar(**50),
      dept_name   **varchar**(20),
      credits   **numeric**(2,0),
      foreign key *(dept_name)* **references** *department)*);

- Primary key declaration can be combined with attribute declaration as shown above

Drop and Alter Table Constructs

- **drop table** *student*
  - Deletes the table and its contents
- **alter table**
  - **alter table** *r* **add** *A D*
    - where *A* is the name of the attribute to be added to relation *r* and *D* is the domain of *A.*
    - All tuples in the relation are assigned *null* as the value for the new attribute.
  - **alter table** *r* **drop** *A*
    - where *A* is the name of an attribute of relation *r*
    - Dropping of attributes not supported by many databases
  - And more …

Basic Query Structure

■ The SQL **data-manipulation language (DML)** provides the ability to query information, and insert, delete and update tuples

■ A typical SQL **query** has the form:

$$\textbf{select } A_1, A_2, ..., A_n$$
$$\textbf{from } r_1, r_2, ..., r_m$$
$$\textbf{where } P$$

● A_i represents an attribute

● R_i represents a relation

● P is a predicate.

■ The result of an SQL query is a **relation**.

The select Clause

- The **select** clause list the attributes desired in the result of a query
  - corresponds to the projection operation of the relational algebra

- Example: find the names of all instructors:

$$\textbf{select } name$$
$$\textbf{from } instructor$$

- NOTE: SQL keywords are case insensitive (i.e., you may use upper- or lower-case letters.)

  - E.g.   *Name* ≡ *NAME* ≡ *name*
  - Some people use upper case wherever we use bold font.

The select Clause (Cont.)

- SQL allows duplicates in relations as well as in query results.

- To force the elimination of duplicates, insert the keyword **distinct** after select.

- Find the names of all departments with instructor, and remove duplicates

    select distinct *dept_name*
    from *instructor*

- The (redundant) keyword **all** specifies that duplicates not be removed.

    select all *dept_name*
    from *instructor*

The select Clause (Cont.)

- An asterisk in the select clause denotes "all attributes"

  select *
  from *instructor*

- The **select** clause can contain arithmetic expressions involving the operation, +, −, *, and /, and operating on constants or attributes of tuples.

  - Most systems also support additional functions
    - ▸ E.g., substring
  - Most systems allow user defined functions (UDFs)

- The query:

  select *ID, name, salary/12*
  from *instructor*

  would return a relation that is the same as the *instructor* relation, except that the value of the attribute *salary* is divided by 12.

The from Clause

- The **from** clause lists the relations involved in the query
  - Corresponds to the Cartesian product operation of the relational algebra.
- Find the Cartesian product *instructor X teaches*

    select *
    from *instructor, teaches*

  - generates every possible instructor – teaches pair, with all attributes from both relations
- Cartesian product not very useful directly, but useful combined with where-clause condition (selection operation in relational algebra)

The where Clause

- The **where** clause specifies conditions that the result must satisfy

    - Corresponds to the selection predicate of the relational algebra.

- To find all instructors in Comp. Sci. dept with salary > 80000

    > **select** *name*
    > **from** *instructor*
    > **where** *dept_name =* 'Comp. Sci.' **and** *salary* > 80000

- Comparison results can be combined using the logical connectives **and, or,** and **not.**

- Comparisons can be applied to results of arithmetic expressions.

- SQL standard: any valid expression that returns a boolean result

    - Vendor specific restrictions may apply!

Cartesian Product: *instructor X teaches*

instructor

ID	name	dept_name	salary
10101	Srinivasan	Comp. Sci.	65000
12121	Wu	Finance	90000
15151	Mozart	Music	40000
22222	Einstein	Physics	95000
32343	El Said	History	60000
33456	Gold	Physics	87000

teaches

ID	course_id	sec_id	semester	year
10101	CS-101	1	Fall	2009
10101	CS-315	1	Spring	2010
10101	CS-347	1	Fall	2009
12121	FIN-201	1	Spring	2010
15151	MU-199	1	Spring	2010
22222	PHY-101	1	Fall	2009

inst.ID	name	dept_name	salary	teaches.ID	course_id	sec_id	semester	year
10101	Srinivasan	Comp. Sci.	65000	10101	CS-101	1	Fall	2009
10101	Srinivasan	Comp. Sci.	65000	10101	CS-315	1	Spring	2010
10101	Srinivasan	Comp. Sci.	65000	10101	CS-347	1	Fall	2009
10101	Srinivasan	Comp. Sci.	65000	12121	FIN-201	1	Spring	2010
10101	Srinivasan	Comp. Sci.	65000	15151	MU-199	1	Spring	2010
10101	Srinivasan	Comp. Sci.	65000	22222	PHY-101	1	Fall	2009
...	...	...	...	...	...	...	...	...
...	...	...	...	...	...	...	...	...
12121	Wu	Finance	90000	10101	CS-101	1	Fall	2009
12121	Wu	Finance	90000	10101	CS-315	1	Spring	2010
12121	Wu	Finance	90000	10101	CS-347	1	Fall	2009
12121	Wu	Finance	90000	12121	FIN-201	1	Spring	2010
12121	Wu	Finance	90000	15151	MU-199	1	Spring	2010
12121	Wu	Finance	90000	22222	PHY-101	1	Fall	2009
...	...	...	...	...	...	...	...	...
...	...	...	...	...	...	...	...	...

Joins

- For all instructors who have taught some course, find their names and the course ID of the courses they taught.

   select *name, course_id*
   from *instructor, teaches*
   where *instructor.ID = teaches.ID*

- Find the course ID, semester, year and title of each course offered by the Comp. Sci. department

   select *section.course_id, semester, year, title*
   from *section, course*
   where *section.course_id = course.course_id* **and**
   dept_name = 'Comp. Sci.'

section
course_id
sec_id
semester
year
building
room_no
time_slot_id

course
course_id
title
dept_name
credits

Try Writing Some Queries in SQL

■ Suggest queries to be written…..

Joined Relations

- **Join operations** take two relations and return as a result another relation.

- A join operation is a Cartesian product which requires that tuples in the two relations match (under some condition). It also specifies the attributes that are present in the result of the join

- The join operations are typically used as subquery expressions in the **from** clause

Join operations – Example

■ Relation *course*

course_id	title	dept_name	credits
BIO-301	Genetics	Biology	4
CS-190	Game Design	Comp. Sci.	4
CS-315	Robotics	Comp. Sci.	3

■ Relation *prereq*

course_id	prereq_id
BIO-301	BIO-101
CS-190	CS-101
CS-347	CS-101

■ Observe that

   prereq information is missing for CS-315 and

   course information is missing for CS-437

Natural Join

- Natural join matches tuples with the same values for all common attributes, and retains only one copy of each common column

  - This is the natural join from relational algebra

- **select** *
  from *instructor* **natural join** *teaches*;

ID	name	dept_name	salary	course_id	sec_id	semester	year
10101	Srinivasan	Comp. Sci.	65000	CS-101	1	Fall	2009
10101	Srinivasan	Comp. Sci.	65000	CS-315	1	Spring	2010
10101	Srinivasan	Comp. Sci.	65000	CS-347	1	Fall	2009
12121	Wu	Finance	90000	FIN-201	1	Spring	2010
15151	Mozart	Music	40000	MU-199	1	Spring	2010
22222	Einstein	Physics	95000	PHY-101	1	Fall	2009
32343	El Said	History	60000	HIS-351	1	Spring	2010
45565	Katz	Comp. Sci.	75000	CS-101	1	Spring	2010
45565	Katz	Comp. Sci.	75000	CS-319	1	Spring	2010
76766	Crick	Biology	72000	BIO-101	1	Summer	2009
76766	Crick	Biology	72000	BIO-301	1	Summer	2010

Natural Join Example

■ List the names of instructors along with the course ID of the courses that they taught.

- **select** *name*, *course_id*
  from *instructor, teaches*
  where *instructor.ID = teaches.ID*;

- **select** *name*, *course_id*
  from *instructor* **natural join** *teaches*;

Natural Join (Cont.)

■ Danger in natural join: beware of unrelated attributes with same name which get equated incorrectly

■ List the names of instructors along with the the titles of courses that they teach

- Incorrect version (makes course.dept_name = instructor.dept_name)
  ‣ **select** *name*, *title*
    from *instructor* **natural join** *teaches* **natural join** *course*;

- Correct version
  ‣ **select** *name*, *title*
    from *instructor* **natural join** *teaches*, *course*
    where *teaches*.*course_id* = *course*.*course_id*;

- Another correct version
  ‣ **select** *name*, *title*
    from (*instructor* **natural join** *teaches*)
                                join *course* **using**(*course_id*);

Outer Join

- An extension of the join operation that avoids loss of information.

- Computes the join and then adds tuples form one relation that does not match tuples in the other relation to the result of the join.

- Uses *null* values.

Left Outer Join

- *course* **natural left outer join** *prereq*

course_id	title	dept_name	credits	prereq_id
BIO-301	Genetics	Biology	4	BIO-101
CS-190	Game Design	Comp. Sci.	4	CS-101
CS-315	Robotics	Comp. Sci.	3	*null*

Right Outer Join

- *course* **natural right outer join** *prereq*

course_id	title	dept_name	credits	prereq_id
BIO-301	Genetics	Biology	4	BIO-101
CS-190	Game Design	Comp. Sci.	4	CS-101
CS-347	*null*	*null*	*null*	CS-101

Joined Relations

- **Join operations** take two relations and return as a result another relation.

- These additional operations are typically used as subquery expressions in the **from** clause

- **Join condition** – defines which tuples in the two relations match, and what attributes are present in the result of the join.

- **Join type** – defines how tuples in each relation that do not match any tuple in the other relation (based on the join condition) are treated.

Join types
inner join
left outer join
right outer join
full outer join

Join Conditions
natural
on <predicate>
using $(A_1, A_1, \ldots, A_n)$

Full Outer Join

- *course* **natural full outer join** *prereq*

course_id	title	dept_name	credits	prereq_id
BIO-301	Genetics	Biology	4	BIO-101
CS-190	Game Design	Comp. Sci.	4	CS-101
CS-315	Robotics	Comp. Sci.	3	*null*
CS-347	*null*	*null*	*null*	CS-101

Joined Relations – Examples

■ *course* **inner join** *prereq* **on**
course.course_id = prereq.course_id

course_id	title	dept_name	credits	prereq_id	course_id
BIO-301	Genetics	Biology	4	BIO-101	BIO-301
CS-190	Game Design	Comp. Sci.	4	CS-101	CS-190

■ What is the difference between the above, and a natural join?

■ *course* **left outer join** *prereq* **on**
course.course_id = prereq.course_id

course_id	title	dept_name	credits	prereq_id	course_id
BIO-301	Genetics	Biology	4	BIO-101	BIO-301
CS-190	Game Design	Comp. Sci.	4	CS-101	CS-190
CS-315	Robotics	Comp. Sci.	3	null	null

Joined Relations – Examples

- *course* **natural right outer join** *prereq*

course_id	title	dept_name	credits	prereq_id
BIO-301	Genetics	Biology	4	BIO-101
CS-190	Game Design	Comp. Sci.	4	CS-101
CS-347	null	null	null	CS-101

- *course* **full outer join** *prereq* **using** (*course_id*)

course_id	title	dept_name	credits	prereq_id
BIO-301	Genetics	Biology	4	BIO-101
CS-190	Game Design	Comp. Sci.	4	CS-101
CS-315	Robotics	Comp. Sci.	3	null
CS-347	null	null	null	CS-101

The Rename Operation

- The SQL allows renaming relations and attributes using the **as** clause:

  old-name **as** *new-name*

- E.g.

  - **select** *ID, name, salary/12* **as** *monthly_salary*
    from *instructor*

- Find the names of all instructors who have a higher salary than some instructor in 'Comp. Sci'.

  - **select distinct** *T. name*
    from *instructor* **as** *T, instructor* **as** *S*
    where *T.salary > S.salary* **and** *S.dept_name =* 'Comp. Sci.'

- Keyword **as** is optional and may be omitted

  instructor **as** *T ≡ instructor T*

  - Keyword **as** must be omitted in Oracle

String Operations

- SQL includes a string-matching operator for comparisons on character strings. The operator "like" uses patterns that are described using two special characters:

    - percent (%). The % character matches any substring.

    - underscore (_). The _ character matches any character.

- Find the names of all instructors whose name includes the substring "dar".

    select *name*
    from *instructor*
    where *name* **like** '%dar%'

- Match the string "100 %"

    like '100 \%' **escape** '\'

String Operations (Cont.)

- Patters are case sensitive.

- Pattern matching examples:
  - 'Intro%' matches any string beginning with "Intro".
  - '%Comp%' matches any string containing "Comp" as a substring.
  - '_ _ _' matches any string of exactly three characters.
  - '_ _ _%' matches any string of at least three characters.

- SQL supports a variety of string operations such as
  - concatenation (using "II")
  - converting from upper to lower case (and vice versa)
  - finding string length, extracting substrings, etc.

Case Construct

- Like case, if, and ? Operators in programming languages

case

> **when** c_1 **then** e_1

> **when** c_2 **then** e_2

> …

> [**else** e_n]

end

- Each c_i is a condition
- Each e_1 is an expression
- Returns the first e_i for which c_i evaluates to *true*
  - If none of the c_i is true, then return e_n (**else**)
    - If there is no else return *null*

Case Construct Example

■ Like case, if, and ? Operators in programming languages

select

    name,

    case

        when salary > 1000000 **then** 'premium'

        else 'standard'

    end as customer_group

from customer

Ordering the Display of Tuples

- List in alphabetic order the names of all instructors

  select distinct *name*
  from *instructor*
  order by *name*

- We may specify **desc** for descending order or **asc** for ascending order, for each attribute; ascending order is the default.

  - Example: **order by** *name* **desc**

- Can sort on multiple attributes

  - Example: **order by** *dept_name, name*

- **Order is not expressible in the relational model!**

Where Clause Predicates

- SQL includes a **between** comparison operator

- Example: Find the names of all instructors with salary between $90,000 and $100,000 (that is, ≥ $90,000 and ≤ $100,000)

  - **select** *name*
    from *instructor*
    where *salary* **between** 90000 **and** 100000

- Tuple comparison

  - **select** *name*, *course_id*
    from *instructor*, *teaches*
    where (*instructor*.*ID*, *dept_name*) = (*teaches*.*ID*, ' Biology');

Set Operations

- Find courses that ran in Fall 2009 or in Spring 2010

  (**select** *course_id* **from** *section* **where** *sem* = 'Fall' **and** *year* = 2009)
   union
  (**select** *course_id* **from** *section* **where** *sem* = 'Spring' **and** *year* = 2010)

- Find courses that ran in Fall 2009 and in Spring 2010

  (**select** *course_id* **from** *section* **where** *sem* = 'Fall' **and** *year* = 2009)
   intersect
  (**select** *course_id* **from** *section* **where** *sem* = 'Spring' **and** *year* = 2010)

- Find courses that ran in Fall 2009 but not in Spring 2010

  (**select** *course_id* **from** *section* **where** *sem* = 'Fall' **and** *year* = 2009)
   except
  (**select** *course_id* **from** *section* **where** *sem* = 'Spring' **and** *year* = 2010)

Set Operations

- Set operations **union**, **intersect**, and **except**

  - Each of the above operations automatically eliminates duplicates

- To retain all duplicates use the corresponding multiset versions **union all, intersect all** and **except all**.

  Suppose a tuple occurs m times in r and n times in s, then, it occurs:

  - $m + n$ times in r **union all** s

  - $\min(m,n)$ times in r **intersect all** s

  - $\max(0, m - n)$ times in r **except all** s

Null Values

- It is possible for tuples to have a null value, denoted by *null*, for some of their attributes

- *null* signifies an unkown value or that a value does not exist.

- The result of any arithmetic expression and comparisons involving *null* evaluate to *null*

  - Example: 5 + *null* returns *null*

    null > 5 returns *null*

    null = *null* returns *null*

- The predicate **is null** can be used to check for null values.

  - Example: Find all instructors whose salary is null*.*

    select *name*
    from *instructor*
    where *salary* **is null**

Null Values and Three Valued Logic

- Any comparison with *null* returns *null*

  - Example*: 5 < null   or   null <> null   or   null = null*

- Three-valued logic using the truth value *null*:

  - OR: (*null* **or** *true*)   = *true*,
        (*null* **or** *false*) = *null*
        (*null* **or** *null*) = *null*

  - AND: *(true* **and** *null)* = *null,*
          (false **and** *null) = false,*
          (null **and** *null) = null*

  - NOT*: (***not*** *null) = null*

  - "*P* **is null**" evaluates to true if predicate *P* evaluates to *null*

- Result of **where** clause predicate is treated as *false* if it evaluates to *null*

Aggregate Functions

■ These functions operate on the multiset of values of a column of a relation, and return a value

avg: average value
min: minimum value
max: maximum value
sum: sum of values
count: number of values

■ Most DBMS support user defined aggregation functions

Aggregate Functions (Cont.)

- Find the average salary of instructors in the Computer Science department

  - **select avg** (*salary*)
    from *instructor*
    where *dept_name*= ' Comp. Sci.' ;

- Find the total number of instructors who teach a course in the Spring 2010 semester

  - **select count** (**distinct** *ID*)
    from *teaches*
    where *semester* = ' Spring' **and** *year* = 2010

- Find the number of tuples in the *course* relation

  - **select count** (*)
    from *course*;

Aggregate Functions – Group By

- Find the average salary of instructors in each department
  - **select** *dept_name*, **avg** (*salary*)
    from *instructor*
    group by *dept_name*;
  - Note: departments with no instructor will not appear in result

ID	name	dept_name	salary
76766	Crick	Biology	72000
45565	Katz	Comp. Sci.	75000
10101	Srinivasan	Comp. Sci.	65000
83821	Brandt	Comp. Sci.	92000
98345	Kim	Elec. Eng.	80000
12121	Wu	Finance	90000
76543	Singh	Finance	80000
32343	El Said	History	60000
58583	Califieri	History	62000
15151	Mozart	Music	40000
33456	Gold	Physics	87000
22222	Einstein	Physics	95000

dept_name	avg_salary
Biology	72000
Comp. Sci.	77333
Elec. Eng.	80000
Finance	85000
History	61000
Music	40000
Physics	91000

Aggregation (Cont.)

- Attributes in **select** clause outside of aggregate functions must appear in **group by** list

  - /* erroneous query */
    select *dept_name*, *ID*, **avg** (*salary*)
    from *instructor*
    group by *dept_name*;

Aggregate Functions – Having Clause

■ Find the names and average salaries of all departments whose average salary is greater than 42000

> **select** *dept_name*, **avg** (*salary*)
> **from** *instructor*
> **group by** *dept_name*
> **having avg** (*salary*) > 42000;

Note: predicates in the **having** clause are applied after the formation of groups whereas predicates in the **where** clause are applied before forming groups

Null Values and Aggregates

- Total all salaries

$$\textbf{select sum }(\textit{salary }) \\ \textbf{from }\textit{instructor}$$

  - Above statement ignores null amounts

  - Result is *null* if there is no non-null amount

- All aggregate operations except **count(*)** ignore tuples with null values on the aggregated attributes

- What if collection has only null values?

  - count returns 0

  - all other aggregates return null

Empty Relations and Aggregates

- What if the input relation is empty

- Conventions:
  - **sum**: returns *null*
  - **avg**: returns *null*
  - **min**: returns *null*
  - **max**: returns *null*
  - **count**: returns 0

Duplicates

- In relations with duplicates, SQL can define how many copies of tuples appear in the result.

- **Multiset** versions of some of the relational algebra operators – given multiset relations r_1 and r_2:

  1. $\sigma_\theta (r_1)$: If there are c_1 copies of tuple t_1 in r_1, and t_1 satisfies selections σ_θ, then there are c_1 copies of t_1 in $\sigma_\theta (r_1)$.

  2. $\Pi_A (r)$: For each copy of tuple t_1 in r_1, there is a copy of tuple $\Pi_A (t_1)$ in $\Pi_A (r_1)$ where $\Pi_A (t_1)$ denotes the projection of the single tuple t_1.

  3. $r_1 \times r_2$: If there are c_1 copies of tuple t_1 in r_1 and c_2 copies of tuple t_2 in r_2, there are $c_1 \times c_2$ copies of the tuple $t_1. t_2$ in $r_1 \times r_2$

Multiset Relational Algebra

- Pure relational algebra operates on **set-semantics** (no duplicates allowed)

  - e.g. after projection

- Multiset (**bag-semantics**) relational algebra retains duplicates, to match SQL semantics

  - SQL duplicate retention was initially for efficiency, but is now a feature

- Multiset relational algebra defined as follows

  - **selection**: has as many duplicates of a tuple as in the input, if the tuple satisfies the selection

  - **projection**: one tuple per input tuple, even if it is a duplicate

  - **cross product**: If there are m copies of $t1$ in r, and n copies of $t2$ in s, there are m x n copies of $t1.t2$ in r x s

  - Other operators similarly defined

    - E.g. **union**: $m + n$ *copies,* **intersection**: min(m, n) copies **difference**: max($0, m - n$) copies

Duplicates (Cont.)

- Example: Suppose multiset relations r_1 (*A, B*) and r_2 (*C*) are as follows:

$$r_1 = \{(1, a)\ (2, a)\} \qquad r_2 = \{(2), (3), (3)\}$$

- Then $\Pi_B(r_1)$ would be $\{(a), (a)\}$, while $\Pi_B(r_1)$ x r_2 would be

$$\{(a,2), (a,2), (a,3), (a,3), (a,3), (a,3)\}$$

- SQL duplicate semantics:

  select $A_1, A_2, ..., A_n$
  from $r_1, r_2, ..., r_m$
  where P

  is equivalent to the *multiset* version of the expression:

$$\prod_{A_1, A_2, \ldots, A_n} (\sigma_P (r_1 \times r_2 \times \ldots \times r_m))$$

SQL and Relational Algebra

- **select** $A_1, A_2, .. A_n$
  from $r_1, r_2, \ldots, r_m$
  where P

  is equivalent to the following expression in multiset relational algebra

  $$\prod_{A1, .., An} (\sigma_P (r_1 \times r_2 \times .. \times r_m))$$

- **select** $A_1, A_2,$ **sum**(A_3)
  from $r_1, r_2, \ldots, r_m$
  where P
  group by A_1, A_2

  is equivalent to the following expression in multiset relational algebra

  $$_{A1, A2} \mathcal{G}_{\textbf{sum}(A3)} (\sigma_P (r_1 \times r_2 \times .. \times r_m)))$$

SQL and Relational Algebra

■ More generally, the non-aggregated attributes in the **select** clause may be a subset of the **group by** attributes, in which case the equivalence is as follows:

select A_1, **sum**(A_3) AS sumA3
from $r_1, r_2, \ldots, r_m$
where P
group by A_1, A_2

is equivalent to the following expression in multiset relational algebra

$$\prod_{A1,sumA3}(\ _{A1,A2}\, \mathcal{G}\, \textbf{sum}(A3)\ \textbf{as}\ \text{sumA3}(\sigma_P(r_1 \times r_2 \times .. \times r_m)))$$

Subqueries in the From Clause

- SQL allows a subquery expression to be used in the **from** clause

- Find the average instructors' salaries of those departments where the average salary is greater than $42,000.

  > **select** *dept_name*, *avg_salary*
  > **from** (**select** *dept_name*, **avg** (*salary*) **as** *avg_salary*
  >   **from** *instructor*
  >   **group by** *dept_name*)
  > **where** *avg_salary* > 42000;

- Note that we do not need to use the **having** clause

- Another way to write above query

  > **select** *dept_name*, *avg_salary*
  > **from** (**select** *dept_name*, **avg** (*salary*)
  >   **from** *instructor*
  >   **group by** *dept_name*)
  >   **as** *dept_avg* (*dept_name*, *avg_salary*)
  > **where** *avg_salary* > 42000;

Nested Subqueries

- SQL provides a mechanism for the nesting of subqueries.

- A **subquery** is a **select-from-where** expression that is nested within another query.

- A common use of subqueries is to perform tests for set membership, set comparisons, and set cardinality.

Example Query

- Find courses offered in Fall 2009 and in Spring 2010

    select distinct *course_id*
    from *section*
    where *semester* = 'Fall' **and** *year*= 2009 **and**
        course_id **in** (**select** *course_id*
                from *section*
                where *semester* = 'Spring' **and** *year*= 2010);

- Find courses offered in Fall 2009 but not in Spring 2010

    select distinct *course_id*
    from *section*
    where *semester* = 'Fall' **and** *year*= 2009 **and**
        course_id **not in** (**select** *course_id*
                from *section*
                where *semester* = 'Spring' **and** *year*=
    2010);

Example Query

■ Find the total number of (distinct) studentswho have taken course sections taught by the instructor with *ID* 10101

> **select count** (**distinct** *ID*)
> **from** *takes*
> **where** (*course_id*, *sec_id*, *semester*, *year*) **in**
>   (**select** *course_id*, *sec_id*, *semester*, *year*
>   **from** *teaches*
>   **where** *teaches*.*ID*= 10101);

■ Note: Above query can be written in a much simpler manner. The formulation above is simply to illustrate SQL features.

Quantification

■ Find names of instructors with salary greater than that of some (at least one) instructor in the Biology department.

> **select distinct** *T.name*
> **from** *instructor* **as** *T*, *instructor* **as** *S*
> **where** *T.salary* > *S.salary* **and** *S.dept_name* = ' Biology ' ;

■ Same query using > **some** clause

> **select** *name*
> **from** *instructor*
> **where** *salary* > **some** (**select** *salary*
>   **from** *instructor*
>   **where** *dept_name* = ' Biology ');

Definition of Some Clause

■ F <comp> **some** $r \Leftrightarrow \exists \, t \in r$ such that (F <comp> t)
Where <comp> can be: $<, \leq, >, =, \neq$

$(5 < \textbf{some} \begin{array}{|c|} \hline 0 \\ \hline 5 \\ \hline 6 \\ \hline \end{array}) = \text{true}$

(read: 5 < some tuple in the relation)

$(5 < \textbf{some} \begin{array}{|c|} \hline 0 \\ \hline 5 \\ \hline \end{array}) = \text{false}$

$(5 = \textbf{some} \begin{array}{|c|} \hline 0 \\ \hline 5 \\ \hline \end{array}) = \text{true}$

$(5 \neq \textbf{some} \begin{array}{|c|} \hline 0 \\ \hline 5 \\ \hline \end{array}) = \text{true (since } 0 \neq 5)$

$(= \textbf{some}) \equiv \textbf{in}$
However, $(\neq \textbf{some}) \not\equiv \textbf{not in}$

Example Query

■ Find the names of all instructors whose salary is greater than the salary of all instructors in the Biology department.

select *name*
from *instructor*
where *salary* > **all** (**select** *salary*
             from *instructor*
             where *dept_name* = ' Biology');

Definition of all Clause

- $F \text{ <comp> } \textbf{all } r \Leftrightarrow \forall\, t \in r \;(F \text{ <comp> } t)$

$$(5 < \textbf{all}\;\; \boxed{\begin{array}{c} 0 \\ 5 \\ 6 \end{array}}\;\;) = \text{false}$$

$$(5 < \textbf{all}\;\; \boxed{\begin{array}{c} 6 \\ 10 \end{array}}\;\;) = \text{true}$$

$$(5 = \textbf{all}\;\; \boxed{\begin{array}{c} 4 \\ 5 \end{array}}\;\;) = \text{false}$$

$$(5 \neq \textbf{all}\;\; \boxed{\begin{array}{c} 4 \\ 6 \end{array}}\;\;) = \text{true (since } 5 \neq 4 \text{ and } 5 \neq 6)$$

$(\neq \textbf{all}) \equiv \textbf{not in}$
However, $(= \textbf{all}) \not\equiv \textbf{in}$

Test for Empty Relations

- The **exists** construct returns the value **true** if the argument subquery returns a nonempty result.

- **exists** $r \Leftrightarrow r \neq \emptyset$

- **not exists** $r \Leftrightarrow r = \emptyset$

Correlation Variables

- Yet another way of specifying the query "Find all courses taught in both the Fall 2009 semester and in the Spring 2010 semester"

  select *course_id*
  from *section* **as** *S*
  where *semester =* 'Fall' **and** *year=* 2009 **and**
      exists (**select** *
          from *section* **as** *T*
          where *semester =* 'Spring' **and** *year=* 2010
             and *S.course_id= T.course_id*);

- **Correlated subquery**

- **Correlation name** or **correlation variable**

Not Exists

- Find all students who have taken all courses offered in the Biology department.

  select distinct *S.ID*, *S.name*
  from *student* **as** *S*
  where not exists ((**select** *course_id*
       from *course*
       where *dept_name* = ' Biology')
      except
       (**select** *T.course_id*
       from *takes* **as** *T*
       where *S.ID = T.ID*));

- Note that $X - Y = \varnothing \quad \Leftrightarrow \quad X \subseteq Y$

- *Note:* Cannot write this query using = **all** and its variants

Test for Absence of Duplicate Tuples

- The **unique** construct tests whether a subquery has any duplicate tuples in its result.

    - (Evaluates to "true" on an empty set)

- Find all courses that were offered at most once in 2009

    select *T.course_id*
    from *course* **as** *T*
    where unique (**select** *R.course_id*
                    from *section* **as** *R*
                    where *T.course_id*= *R.course_id*
                            and *R.year* = 2009);

Correlated Subqueries in the From Clause

- And yet another way to write it: **lateral** clause

  > **select** *name*, *salary*, *avg_salary*
  > **from** *instructor I1*,
  >   **lateral** (**select avg**(*salary*) as *avg_salary*
  >   **from** *instructor I2*
  >   **where** *I2.dept_name*= *I1.dept_name*);

- Lateral clause permits later part of the **from** clause (after the lateral keyword) to access correlation variables from the earlier part.

- Note: lateral is part of the SQL standard, but is not supported on many database systems; some databases such as SQL Server offer alternative syntax

With Clause

■ The **with** clause provides a way of defining a temporary view whose definition is available only to the query in which the **with** clause occurs.

■ Find all departments with the maximum budget

> **with** *max_budget* (*value*) **as**
>   (**select max**(*budget*)
>   **from** *department*)
> **select** *budget*
> **from** *department, max_budget*
> **where** *department.budget = max_budget.value*;

Complex Queries using With Clause

■ With clause is very useful for writing complex queries

■ Supported by most database systems, with minor syntax variations

■ Find all departments where the total salary is greater than the average of the total salary at all departments

with *dept _total* (*dept_name*, *value*) **as**
     (**select** *dept_name*, **sum**(*salary*)
      from *instructor*
      group by *dept_name*),
dept_total_avg(*value*) **as**
     (**select avg**(*value*)
     from *dept_total*)
select *dept_name*
from *dept_total*, *dept_total_avg*
where *dept_total.value* >= *dept_total_avg.value*;

Scalar Subquery

- Scalar subquery is one which is used where a single value is expected

- E.g. **select** *dept_name,*
        (**select count**(*)
          from *instructor*
         where *department.dept_name = instructor.dept_name*)
        as *num_instructors*
      from *department*;

- E.g. **select** *name*
      from *instructor*
      where *salary * 10 >*
        (**select** *budget* **from** *department*
         where *department.dept_name = instructor.dept_name*)

- Runtime error if subquery returns more than one result tuple

Query Features Recap - Syntax

■ An SQL query is either a Select-from-where block or a set operation

■ An SQL query block is structured like this:

SELECT [**DISTINCT**] select_list

[**FROM** from_list]

[**WHERE** where_condition]

[**GROUP BY** group_by_list]

[**HAVING** having_condition]

[**ORDER BY** order_by_list]

■ Set operations

[Query Block] set_op [Query Block]

set_op: [**ALL**] **UNION** | **INTERSECT** | **EXCEPT**

Query Features Recap - Syntax

- Almost all clauses are optional
- Examples:
  - **SELECT * FROM** r;
  - **SELECT** 1;
    - ▸ Convention: returns single tuple
  - **SELECT** 'ok' **FROM** accounts **HAVING** sum(balance) = 0;
  - **SELECT** 1 **GROUP BY** 1;
  - **SELECT** 1 **HAVING** true;
  - Let r be a relation with two attributes a and b
    - ▸ **SELECT** a,b **FROM** r

      WHERE a **IN** (**SELECT** a **FROM** r) **AND** b **IN** (**SELECT** b **FROM** r)

      GROUP BY a,b **HAVING** count(*) > 0;
- Note:
  - Not all systems support all of this "non-sense"

Syntax - SELECT

- **SELECT** [**DISTINCT [ON (distinct_list)]**] select_list

- select_list

  - List of projection expressions

    - [expr] [**AS** name]

  - expr

    - Expression over attributes, constants, arithmetic operators, functions, **CASE**-construct, aggregation functions

- distinct_list

  - List of expressions

- Examples:

  - **SELECT DISTINCT ON** (a % 2) a **FROM** r;

  - **SELECT** substring(a, 1,2) **AS** x **FROM** r;

  - **SELECT CASE WHEN** a = 2 **THEN** a **ELSE** null **END AS** b **FROM** r;

  - **SELECT** a = b **AS** is_a_equal_to_b **FROM** r;

Syntax - FROM

- **FROM** from_list

- from_list
  - List of from clause expressions
    - subquery | relation | constant_relation | join_expr [alias]
  - subquery
    - Any valid SQL query – alias is not optional
  - relation
    - A relation in the database
  - constant_relation
    - (**VALUES** tuples) – alias is not optional
  - join_expr
    - joins between from_clause entries
  - alias
    - [AS] b [(attribute_name_list)]

Syntax – FROM (cont.)

■ Examples (relation r with attributes a and b):

- **SELECT** * **FROM** r;

- **SELECT** * **FROM** r **AS** g(v,w);

- **SELECT** * **FROM** r x;

- **SELECT** * **FROM** (**VALUES** (1,2), (3,1)) **AS** s(u,v);

- **SELECT** * **FROM** r **NATURAL JOIN** s, t;

- **SELECT** * **FROM** ((r **JOIN** s **ON** (r.a = s.c)) **NATURAL JOIN** (**SELECT** * **FROM** t) **AS** new);

- **SELECT** * **FROM** (**SELECT** * **FROM** r) **AS** r;

- **SELECT** * **FROM** (**SELECT** * **FROM** (**SELECT** * **FROM** r) **AS** r) **AS** r;

Syntax - WHERE

- **WHERE** where_condition

- where_condition: A boolean expression over

  - Attributes

  - Constants: e.g., true, 1, 0.5, 'hello'

  - Comparison operators: =, <, >, **IS DISTINCT FROM, IS NULL**, …

  - Arithmetic operators: +,-,/,%

  - Function calls

  - Nested subquery expressions

- Examples

  - **SELECT** * **FROM** r **WHERE** a = 2;

  - **SELECT** * **FROM** r **WHERE** true **OR** false**;**

  - **SELECT** * **FROM** r **WHERE NOT**(a = 2 **OR** a = 3)**;**

  - **SELECT** * **FROM** r **WHERE** a **IS DISTINCT FROM** b**;**

  - **SELECT** * **FROM** r **WHERE** a < **ANY** (**SELECT** c **FROM** s)**;**

  - **SELECT** * **FROM** r **WHERE** a = (**SELECT** count(*) **FROM** s);

Syntax – GROUP BY

- **GROUP BY** group_by_list
- group_by_list
  - List of expressions
    - ▸ Expression over attributes, constants, arithmetic operators, functions, **CASE**-construct, aggregation functions
- Examples:
  - **SELECT** sum(a), b **FROM** r **GROUP BY** b;
  - **SELECT** sum(a), b, c **FROM** r **GROUP BY** b, c;
  - **SELECT** sum(a), b/2 **FROM** r **GROUP BY** b/2;
  - **SELECT** sum(a), b **FROM** r **GROUP BY** b > 5;
    - ▸ Incorrect, cannot select b, because it is not an expression in the group by clause
  - **SELECT** sum(a), b **FROM** r **GROUP BY** b **IN** (**SELECT** c **FROM** s);

Syntax – HAVING

- **HAVING** having_condition
- having_condition
  - Like where_condition except that expressions over attributes have either to be in the **GROUP BY** clause or are aggregated
- Examples:
  - **SELECT** sum(a), b **FROM** r **GROUP BY** b **HAVING** sum(a) > 10;
  - **SELECT** sum(a), b **FROM** r **GROUP BY** b **HAVING** sum(a) + 5 > 10;
  - **SELECT** sum(a), b **FROM** r **GROUP BY** b **HAVING** true;
  - **SELECT** sum(a), b **FROM** r **GROUP BY** b **HAVING** count(*) = 50;
  - **SELECT** b **FROM** r **GROUP BY** b **HAVING** sum(a) > 10;

Syntax – ORDER BY

- **ORDER BY** order_by_list

- order_by_list
  - Like select_list minus renaming
  - Optional [**ASC** | **DESC**] for each item

- Examples:
  - **SELECT * FROM** r **ORDER BY** a;
  - **SELECT * FROM** r **ORDER BY** b, a;
  - **SELECT * FROM** r **ORDER BY** a * 2;
  - **SELECT * FROM** r **ORDER BY** a * 2, a;
  - **SELECT * FROM** r **ORDER BY** a + (SELECT count(*) FROM s);

Query Semantics

- Evaluation Algorithm (you can do it manually – sort of)

1. Compute **FROM** clause

    1. Compute cross product of all items in the **FROM** clause

        - Relations: nothing to do

        - Subqueries: use this algorithm to recursively compute the result of subqueries first

        - Join expressions: compute the join

2. Compute **WHERE** clause

    1. For each tuple in the result of 1. evaluate the **WHERE** clause condition

3. Compute **GROUP BY** clause

    1. Group the results of step 2. on the **GROUP BY** expressions

4. Compute **HAVING** clause

    1. For each group (if any) evaluate the **HAVING** condition

Query Semantics (Cont.)

5. Compute **ORDER BY** clause

    5. Order the result of step 4 on the **ORDER BY** expressions

6. Compute **SELECT** clause

    5. Project each result tuple from step 5 on the **SELECT** expressions

- If the **WHERE**, **SELECT**, **GROUP BY**, **HAVING**, **ORDER BY** clauses have any nested subqueries
  - For each tuple t in the result of the **FROM** clause
    - Substitute the correlated attributes with values from t
    - Evaluate the resulting query
    - Use the result to evaluate the expression in the clause the subquery occurs in

Query Semantics (Cont.)

- Equivalent relational algebra expression

    - **ORDER BY** has no equivalent, because relations are unordered

    - Nested subqueries: need to extend algebra (not covered here)

- Each query block is equivalent to

$$\pi(\sigma(\mathcal{G}(\pi(\sigma(F_1 \times \ldots F_n)))))$$

- Where F_i is the translation of the i^{th} **FROM** clause item

- Note: we leave out the arguments

Modification of the Database

- Deletion of tuples from a given relation

- Insertion of new tuples into a given relation

- Updating values in some tuples in a given relation

Modification of the Database – Deletion

■ Delete all instructors

        delete from *instructor*

■ Delete all instructors from the Finance department
        delete from *instructor*
        where *dept_name*= ' Finance' ;

■ Delete all tuples in the *instructor* relation for those instructors associated with a department located in the Watson building.

        delete from *instructor*
        where *dept_name* **in** (**select** *dept_name*
                  from *department*
                  where *building* = ' Watson');

Deletion (Cont.)

■ Delete all instructors whose salary is less than the average salary of instructors

delete from *instructor*
where *salary* < (**select avg** (*salary*) **from** *instructor*);

● Problem: as we delete tuples from instructor, the average salary changes

● Solution used in SQL:

1. First, compute **avg** salary and find all tuples to delete

2. Next, delete all tuples found above (without recomputing **avg** or retesting the tuples)

Modification of the Database – Insertion

■ Add a new tuple to *course*

> **insert into** *course*
> > **values** (' CS-437' , ' Database Systems' , ' Comp. Sci.' , 4);

■ or equivalently

> **insert into** *course* (*course_id*, *title*, *dept_name*, *credits*)
> > **values** (' CS-437' , ' Database Systems' , ' Comp. Sci.' , 4);

■ Add a new tuple to *student* with *tot_creds* set to null

> **insert into** *student*
> > **values** (' 3003' , ' Green' , ' Finance' , *null*);

Insertion (Cont.)

■ Add all instructors to the *student* relation with tot_creds set to 0

> **insert into** *student*
>   **select** *ID, name, dept_name, 0*
>   **from** *instructor*

■ The **select from where** statement is evaluated fully before any of its results are inserted into the relation (otherwise queries like

> **insert into** *table*1 **select** * **from** *table*1

would cause problems, if *table1* did not have any primary key defined.

Modification of the Database – Updates

- Increase salaries of instructors whose salary is over $100,000 by 3%, and all others receive a 5% raise

  - Write two **update** statements:

    > **update** *instructor*
    >   **set** *salary* = *salary* * 1.03
    >   **where** *salary* > 100000;
    > **update** *instructor*
    >   **set** *salary* = *salary* * 1.05
    >   **where** *salary* <= 100000;

  - The order is important

  - Can be done better using the **case** statement (next slide)

Case Statement for Conditional Updates

- Same query as before but with case statement

    update *instructor*
      set *salary* = **case**
                    when *salary* <= 100000 **then** *salary* * 1.05
                    else *salary* * 1.03
                    end

Updates with Scalar Subqueries

- Recompute and update tot_creds value for all students

  update *student S*
     set *tot_cred* = (**select sum**(*credits*)
                              from *takes* **natural join** *course*
                           where *S.ID= takes.ID* **and**
                                  takes.grade $\Leftrightarrow$ ' F' **and**
                                  takes.grade **is not null**);

- Sets *tot_creds* to null for students who have not taken any course

- Instead of **sum**(*credits*), use:

        case
           when sum(*credits*) **is not null then sum**(*credits*)
           else 0
        end

- Or **COALESCE**(sum(credits),0)
  - **COALESCE** returns first non-null arguments

Recap

- SQL queries
  - Clauses: **SELECT**, **FROM** , **WHERE**, **GROUP BY**, **HAVING**, **ORDER BY**
  - Nested subqueries
  - Equivalence with relational algebra
- SQL update, inserts, deletes
  - Semantics of referencing updated relation in **WHERE**
- SQL DDL
  - Table definition: **CREATE TABLE**

End of Chapter 4

Outline

- Introduction
- Relational Data Model
- Formal Relational Languages (relational algebra)
- **SQL - Intermediate**
- Database Design
- Transaction Processing, Recovery, and Concurrency Control
- Storage and File Structures
- Indexing and Hashing
- Query Processing and Optimization

Advanced SQL Features**

- Create a table with the same schema as an existing table:

  create table *temp_account* **like** *account*

Figure 3.02

name
Srinivasan
Wu
Mozart
Einstein
El Said
Gold
Katz
Califieri
Singh
Crick
Brandt
Kim

Figure 3.03

dept_name
Comp. Sci.
Finance
Music
Physics
History
Physics
Comp. Sci.
History
Finance
Biology
Comp. Sci.
Elec. Eng.

Figure 3.04

name
Katz
Brandt

Figure 3.05

name	dept_name	building
Srinivasan	Comp. Sci.	Taylor
Wu	Finance	Painter
Mozart	Music	Packard
Einstein	Physics	Watson
El Said	History	Painter
Gold	Physics	Watson
Katz	Comp. Sci.	Taylor
Califieri	History	Painter
Singh	Finance	Painter
Crick	Biology	Watson
Brandt	Comp. Sci.	Taylor
Kim	Elec. Eng.	Taylor

Figure 3.07

name	Course_id
Srinivasan	CS-101
Srinivasan	CS-315
Srinivasan	CS-347
Wu	FIN-201
Mozart	MU-199
Einstein	PHY-101
El Said	HIS-351
Katz	CS-101
Katz	CS-319
Crick	BIO-101
Crick	BIO-301
Brandt	CS-190
Brandt	CS-190
Brandt	CS-319
Kim	EE-181

Figure 3.08

ID	name	dept_name	salary	course_id	sec_id	semester	year
10101	Srinivasan	Comp. Sci.	65000	CS-101	1	Fall	2009
10101	Srinivasan	Comp. Sci.	65000	CS-315	1	Spring	2010
10101	Srinivasan	Comp. Sci.	65000	CS-347	1	Fall	2009
12121	Wu	Finance	90000	FIN-201	1	Spring	2010
15151	Mozart	Music	40000	MU-199	1	Spring	2010
22222	Einstein	Physics	95000	PHY-101	1	Fall	2009
32343	El Said	History	60000	HIS-351	1	Spring	2010
45565	Katz	Comp. Sci.	75000	CS-101	1	Spring	2010
45565	Katz	Comp. Sci.	75000	CS-319	1	Spring	2010
76766	Crick	Biology	72000	BIO-101	1	Summer	2009
76766	Crick	Biology	72000	BIO-301	1	Summer	2010
83821	Brandt	Comp. Sci.	92000	CS-190	1	Spring	2009
83821	Brandt	Comp. Sci.	92000	CS-190	2	Spring	2009
83821	Brandt	Comp. Sci.	92000	CS-319	2	Spring	2010
98345	Kim	Elec. Eng.	80000	EE-181	1	Spring	2009

Figure 3.09

course_id
CS-101
CS-347
PHY-101

Figure 3.10

course_id
CS-101
CS-315
CS-319
CS-319
FIN-201
HIS-351
MU-199

Figure 3.11

course_id
CS-101
CS-315
CS-319
CS-347
FIN-201
HIS-351
MU-199
PHY-101

Figure 3.12

course_id
CS-101

Figure 3.13

course_id
CS-347
PHY-101

Figure 3.16

dept_name	count
Comp. Sci.	3
Finance	1
History	1
Music	1

Figure 3.17

dept_name	avg(salary)
Physics	91000
Elec. Eng.	80000
Finance	85000
Comp. Sci.	77333
Biology	72000
History	61000

CS425 – Fall 2013
Boris Glavic
Chapter 5: Intermediate SQL

modified from:

Database System Concepts, 6[th] Ed.

Chapter 5: Intermediate SQL

■ Views

■ Transactions

■ Integrity Constraints

■ SQL Data Types and Schemas

■ Access Control

Textbook: Chapter 4

Views

- In some cases, it is not desirable for all users to see the entire logical model (that is, all the actual relations stored in the database.)

- Consider a person who needs to know an instructors name and department, but not the salary. This person should see a relation described, in SQL, by

> **select** *ID*, *name*, *dept_name*
> **from** *instructor*

- A **view** provides a mechanism to hide certain data from the view of certain users.

- Any relation that is not of the conceptual model but is made visible to a user as a "virtual relation" is called a **view**.

View Definition

- A view is defined using the **create view** statement which has the form

  create view v **as** $<$ query expression $>$

  where <query expression> is any legal SQL expression. The view name is represented by v.

- Once a view is defined, the view name can be used to refer to the virtual relation that the view generates.

- View definition is not the same as creating a new relation by evaluating the query expression

  - Rather, a view definition causes the saving of an expression; the expression is substituted into queries using the view.

Example Views

- A view of instructors without their salary

  create view *faculty* **as**
  > **select** *ID*, *name*, *dept_name*
  > **from** *instructor*

- Find all instructors in the Biology department

  select *name*
  from *faculty*
  where *dept_name* = 'Biology'

- Create a view of department salary totals

  create view *departments_total_salary*(*dept_name*, *total_salary*) **as**
  > **select** *dept_name*, **sum** (*salary*)
  > **from** *instructor*
  > **group by** *dept_name*;

Views Defined Using Other Views

- **create view** *physics_fall_2009* **as**
  select *course.course_id*, *sec_id*, *building*, *room_number*
  from *course, section*
  where *course.course_id = section.course_id*
            and *course.dept_name* = 'Physics'
            and *section.semester* = 'Fall'
            and *section.year* = '2009';

- **create view** *physics_fall_2009_watson* **as**
  select *course_id*, *room_number*
  from *physics_fall_2009*
  where *building*= 'Watson';

View Expansion

- Expand use of a view in a query/another view

create view *physics_fall_2009_watson* **as**
(**select** *course_id*, *room_number*
from (**select** *course.course_id*, *building*, *room_number*
 from *course*, *section*
 where *course.course_id* = *section.course_id*
  and *course.dept_name* = ' Physics'
  and *section.semester* = ' Fall'
  and *section.year* = ' 2009')
where *building*= ' Watson' ;

Views Defined Using Other Views

- One view may be used in the expression defining another view

- A view relation v_1 is said to *depend directly* on a view relation v_2 if v_2 is used in the expression defining v_1

- A view relation v_1 is said to *depend on* view relation v_2 if either v_1 depends directly to v_2 or there is a path of dependencies from v_1 to v_2

- A view relation v is said to be *recursive* if it depends on itself.

View Expansion

- A way to define the meaning of views defined in terms of other views.

- Let view v_1 be defined by an expression e_1 that may itself contain uses of view relations.

- View expansion of an expression repeats the following replacement step:

  repeat
      Find any view relation v_i in e_1
      Replace the view relation v_i by the expression defining v_i
  until no more view relations are present in e_1

- As long as the view definitions are not recursive, this loop will terminate

Update of a View

■ Add a new tuple to *faculty* view which we defined earlier

insert into *faculty* **values** ('30765', 'Green', 'Music');

This insertion must be represented by the insertion of the tuple

('30765', 'Green', 'Music', null)

into the *instructor* relation

Some Updates cannot be Translated Uniquely

- **create view** *instructor_info* **as**
    select *ID*, *name*, *building*
     from *instructor*, *department*
     where *instructor*.*dept_name*= *department*.*dept_name*;

- **insert into** *instructor_info* **values** ('69987', 'White', 'Taylor');

    ▸ which department, if multiple departments in Taylor?

    ▸ what if no department is in Taylor?

- Most SQL implementations allow updates only on simple views

    - The **from** clause has only one database relation.

    - The **select** clause contains only attribute names of the relation, and does not have any expressions, aggregates, or **distinct** specification.

    - Any attribute not listed in the **select** clause can be set to null

    - The query does not have a **group** by or **having** clause.

And Some Not at All

- **create view** *history_instructors* **as**
  select *
  from *instructor*
  where *dept_name*= ꞌHistoryꞌ;

- What happens if we insert (ꞌ25566ꞌ, ꞌBrownꞌ, ꞌBiologyꞌ, 100000) into *history_instructors?*

Materialized Views

■ **Materializing a view**: create a physical table containing all the tuples in the result of the query defining the view

■ If relations used in the query are updated, the materialized view result becomes out of date

  ● Need to **maintain** the view, by updating the view whenever the underlying relations are updated.

Transactions

Transactions

- Unit of work

- Atomic transaction

  - either fully executed or rolled back as if it never occurred

- Isolation from concurrent transactions

- Transactions begin implicitly

  - Ended by **commit work** or **rollback work**

- But default on most databases: each SQL statement commits automatically

  - Can turn off auto commit for a session (e.g. using API)

  - In SQL:1999, can use: **begin atomic** …. **end**

    ▸ Not supported on most databases

Transactions Example

- Example Atomicity (all-or-nothing)

  - Recall example from the introduction

  - Relation `accounts(accID, cust, type, balance)`

  - A user want to transfer $100 from his savings (accID = *100*) to his checking account (*accID= 101*)

  UPDATE accounts **SET** balance = balance – 100 **WHERE** accID = 100;

  UPDATE accounts **SET** balance = balance + 100 **WHERE** accID = 101;

  - This can cause inconsistencies if the system crashes after the first update (user would loose money)

  - Using a transaction either both or none of the statements are executed

BEGIN

  UPDATE accounts **SET** balance = balance – 100 **WHERE** accID = 100;

  UPDATE accounts **SET** balance = balance + 100 **WHERE** accID = 101;

COMMIT

Transactions and Concurrency

- Transactions are also used to isolate concurrent actions of different users

- Recall from the introduction that if several users are modifying the database at the same time that can lead to inconsistencies

- More on that later once we talk about concurrency control

Integrity Constraints

modified from:

Database System Concepts, 6th Ed.

Integrity Constraints

- Integrity constraints guard against accidental damage to the database, by ensuring that authorized changes to the database do not result in a loss of data consistency.

  - A checking account must have a balance greater than $10,000.00

  - A salary of a bank employee must be at least $4.00 an hour

  - A customer must have a (non-null) phone number

Integrity Constraints on a Single Relation

- **not null**

- **primary key**

- **unique**

- **check** (P), where P is a predicate

Not Null and Unique Constraints

- **not null**

  - Declare *name* and *budget* to be **not null**

    name **varchar**(20) **not null**
    budget **numeric**(12,2) **not null**

- **unique** (A_1, A_2, …, A_m)

  - The unique specification states that the attributes $A1$, $A2$, … Am
    form a candidate key.

  - Candidate keys are permitted to be null (in contrast to primary keys).

The check clause

■ **check** (P)

where P is a predicate

Example: ensure that semester is one of fall, winter, spring or summer:

create table *section* (
    course_id **varchar** (8),
    sec_id **varchar** (8),
    semester **varchar** (6),
    year **numeric** (4,0),
    building **varchar** (15),
    room_number **varchar** (7),
    time slot id **varchar** (4),
    primary key (*course_id*, *sec_id*, *semester*, *year*),
    check (*semester* **in** (' Fall' , ' Winter' , ' Spring' ,
' Summer'))
);

Referential Integrity

- Ensures that a value that appears in one relation for a given set of attributes also appears for a certain set of attributes in another relation.

  - Example: If "Biology" is a department name appearing in one of the tuples in the *instructor* relation, then there exists a tuple in the *department* relation for "Biology".

- Let A be a set of attributes. Let R and S be two relations that contain attributes A and where A is the primary key of S. A is said to be a **foreign key** of R if for any values of A appearing in R these values also appear in S.

Cascading Actions in Referential Integrity

- **create table** *course* (
  course_id   **char**(5) **primary key**,
  title   **varchar**(20),
  dept_name **varchar**(20) **references** *department*
)

- **create table** *course* (

  …
  dept_name **varchar**(20),
  foreign key (*dept_name*) **references** *department*
         on delete cascade
         on update cascade,
  . . .
)

- alternative actions to cascade: **set null**, **set default**

Integrity Constraint Violation During Transactions

■ E.g.

  create table *person* (
    ID **char**(10),
    name **char**(40),
    mother **char**(10),
    father **char**(10),
    primary key *ID,*
    foreign key *father* **references** *person,*
    foreign key *mother* **references** *person*)

■ How to insert a tuple without causing constraint violation ?

  ● insert father and mother of a person before inserting person

  ● OR, set father and mother to null initially, update after inserting all persons (not possible if father and mother attributes declared to be **not null**)

  ● OR defer constraint checking (next slide)

Complex Check Clauses

- **check** (*time_slot_id* **in** (**select** *time_slot_id* **from** *time_slot*))
  - why not use a foreign key here?
- Every section has at least one instructor teaching the section.
  - how to write this?
- Unfortunately: subquery in check clause not supported by pretty much any database
  - Alternative: triggers (later)
- **create assertion** <assertion-name> **check** <predicate>;
  - Also not supported by anyone

Indexes and User-Defined Types (UDTs)

modified from:

Database System Concepts, 6th Ed.

Built-in Data Types in SQL

- **date:** Dates, containing a (4 digit) year, month and date
  - Example: **date** '2005-7-27'
- **time:** Time of day, in hours, minutes and seconds.
  - Example: **time** '09:00:30'   **time** '09:00:30.75'
- **timestamp:** date plus time of day
  - Example: **timestamp** '2005-7-27 09:00:30.75'
- **interval:** period of time
  - Example: interval '1' day
  - Subtracting a date/time/timestamp value from another gives an interval value
  - Interval values can be added to date/time/timestamp values

Index Creation

- **create table** *student*
  (*ID* **varchar** (5),
  name **varchar** (20) **not null**,
  dept_name **varchar** (20),
  tot_cred **numeric** (3,0) **default** 0,
  primary key (*ID*))

- **create index** *studentID_index* **on** *student*(*ID*)

- Indices are data structures used to speed up access to records with specified values for index attributes

  - e.g. **select** *
    from *student*
    where *ID* = '12345'

  can be executed by using the index to find the required record, without looking at all records of *student*

  More on indices later

User-Defined Types

■ **create type** construct in SQL creates user-defined type

        create type *Dollars* **as numeric (12,2) final**

  ● **create table** *department*
    (*dept_name* **varchar** (20),
    building **varchar** (15),
    budget Dollars);

Domains

- **create domain** construct in SQL-92 creates user-defined domain types

  create domain *person_name* **char**(20) **not null**

- Types and domains are similar. Domains can have constraints, such as **not null**, specified on them.

- **create domain** *degree_level* **varchar**(10)
  constraint *degree_level_test*
  check (**value in** ('Bachelors', 'Masters', 'Doctorate'));

Large-Object Types

- Large objects (photos, videos, CAD files, etc.) are stored as a *large object*:

  - **blob**: binary large object -- object is a large collection of uninterpreted binary data (whose interpretation is left to an application outside of the database system)

  - **clob**: character large object -- object is a large collection of character data

  - When a query returns a large object, a pointer is returned rather than the large object itself.

Access Control

modified from:

Database System Concepts, 6ᵗʰ Ed.

©Silberschatz, Korth and Sudarshan
See www.db-book.com for conditions on re-use

Access Control

Forms of authorization on parts of the database:

- **Read** - allows reading, but not modification of data.

- **Insert** - allows insertion of new data, but not modification of existing data.

- **Update** - allows modification, but not deletion of data.

- **Delete** - allows deletion of data.

Forms of authorization to modify the database schema

- **Index** - allows creation and deletion of indices.

- **Resources** - allows creation of new relations.

- **Alteration** - allows addition or deletion of attributes in a relation.

- **Drop** - allows deletion of relations.

Authorization Specification in SQL

■ The **grant** statement is used to confer authorization

    grant <privilege list>

    on <relation name or view name> **to** <user list>

■ <user list> is:

  ● a user-id

  ● **public**, which allows all valid users the privilege granted

  ● A role (more on this later)

■ Granting a privilege on a view does not imply granting any privileges on the underlying relations.

■ The grantor of the privilege must already hold the privilege on the specified item (or be the database administrator).

Privileges in SQL

■ **select:** allows read access to relation, or the ability to query using the view

  ● Example: grant users U_1, U_2, and U_3 **select** authorization on the *instructor* relation:

  > **grant select on** *instructor* **to** U_1, U_2, U_3

■ **insert**: the ability to insert tuples

■ **update**: the ability to update using the SQL update statement

■ **delete**: the ability to delete tuples.

■ **all privileges**: used as a short form for all the allowable privileges

Revoking Authorization in SQL

- The **revoke** statement is used to revoke authorization.

  revoke <privilege list>

  on <relation name or view name> **from** <user list>

- Example:

  revoke select on *branch* **from** U_1, U_2, U_3

- <privilege-list> may be **all** to revoke all privileges the revokee may hold.

- If <revokee-list> includes **public,** all users lose the privilege except those granted it explicitly.

- If the same privilege was granted twice to the same user by different grantees, the user may retain the privilege after the revocation.

- All privileges that depend on the privilege being revoked are also revoked.

Roles

- **create role** instructor;

- **grant** *instructor* **to Amit;**

- Privileges can be granted to roles:

  - **grant select on** *takes* **to** *instructor*;

- Roles can be granted to users, as well as to other roles

  - **create role** *teaching_assistant*

  - **grant** *teaching_assistant* **to** *instructor*;

    - *Instructor* inherits all privileges of *teaching_assistant*

- Chain of roles

  - **create role** *dean*;

  - **grant** *instructor* **to** *dean*;

  - **grant** *dean* **to** Satoshi;

Authorization on Views

- **create view** *geo_instructor* **as**
  (**select** *
  from *instructor*
  where *dept_name* = ' Geology');

- **grant select on** *geo_instructor* **to** *geo_staff*

- Suppose that a *geo_staff* member issues

  - **select** *
    from *geo_instructor*;

- What if

  - *geo_staff* does not have permissions on *instructor?*

  - creator of view did not have some permissions on *instructor?*

Other Authorization Features

- **references** privilege to create foreign key
  - **grant reference** (*dept_name*) **on** *department* **to** Mariano;
  - why is this required?
- transfer of privileges
  - **grant select on** *department* **to** Amit **with grant option**;
  - **revoke select on** *department* **from** Amit, Satoshi **cascade**;
  - **revoke select on** *department* **from** Amit, Satoshi **restrict**;
- Etc. read text book Section 4.6 for more details we have omitted here.

Understanding RESTRICT/CASCADE

- Bob grants right X on Y to Alice with grant option
- Alice grants right X on Y to Peter

- **Abandoned right**
  - A right for which there is no justification anymore

- **revoke** X **on** *Y* **from** Peter **restrict**
  - With restrict fails if it would result in abandoned rights

- **revoke** X **on** *Y* **from** Peter **cascade**
  - Also revokes rights that would otherwise be abandoned

Understanding RESTRICT/CASCADE

- Bob grants right X on Y to Alice with grant option
- Alice grants right X on Y to Peter
- Bob grants right X on Y to Peter

- **Abandoned privilege**
  - A privilege for which there is no justification anymore
  - Indirect justifications count
- **revoke** X **on** *Y* **from** Peter **restrict**
  - Fails: even though there exists additional justification for the privilege.
- **revoke** X **on** *Y* **from** Peter **cascade**
  - Revokes that right from Peter.
  - Peter still has the right to do X on Y

Recap

- Views
  - Virtual
  - Materialized
  - Updates
- Integrity Constraints
  - Not null, unique, check
  - Foreign keys: referential integrity
- Access control
  - Users, roles
  - Privileges
  - **GRANT** / **REVOKE**
- Data types
  - Build-in types, Domains, Large Objects
  - UDTs
  - Indices

End of Chapter 5

Outline

- Introduction
- Relational Data Model
- Formal Relational Languages (relational algebra)
- **SQL - Advanced**
- Database Design
- Transaction Processing, Recovery, and Concurrency Control
- Storage and File Structures
- Indexing and Hashing
- Query Processing and Optimization

Figure 4.01

ID	name	dept_name	tot_cred
00128	Zhang	Comp. Sci.	102
12345	Shankar	Comp. Sci.	32
19991	Brandt	History	80
23121	Chavez	Finance	110
44553	Peltier	Physics	56
45678	Levy	Physics	46
54321	Williams	Comp. Sci.	54
55739	Sanchez	Music	38
70557	Snow	Physics	0
76543	Brown	Comp. Sci.	58
76653	Aoi	Elec. Eng.	60
98765	Bourikas	Elec. Eng.	98
98988	Tanaka	Biology	120

Figure 4.02

ID	course_id	sec_id	semester	year	grade
00128	CS-101	1	Fall	2009	A
00128	CS-347	1	Fall	2009	A-
12345	CS-101	1	Fall	2009	C
12345	CS-190	2	Spring	2009	A
12345	CS-315	1	Spring	2010	A
12345	CS-347	1	Fall	2009	A
19991	HIS-351	1	Spring	2010	B
23121	FIN-201	1	Spring	2010	C+
44553	PHY-101	1	Fall	2009	B-
45678	CS-101	1	Fall	2009	F
45678	CS-101	1	Spring	2010	B+
45678	CS-319	1	Spring	2010	B
54321	CS-101	1	Fall	2009	A-
54321	CS-190	2	Spring	2009	B+
55739	MU-199	1	Spring	2010	A-
76543	CS-101	1	Fall	2009	A
76543	CS-319	2	Spring	2010	A
76653	EE-181	1	Spring	2009	C
98765	CS-101	1	Fall	2009	C-
98765	CS-315	1	Spring	2010	B
98988	BIO-101	1	Summer	2009	A
98988	BIO-301	1	Summer	2010	null

Figure 4.03

ID	name	dept_name	tot_cred	course_id	sec_id	semester	year	grade
00128	Zhang	Comp. Sci.	102	CS-101	1	Fall	2009	A
00128	Zhang	Comp. Sci.	102	CS-347	1	Fall	2009	A-
12345	Shankar	Comp. Sci.	32	CS-101	1	Fall	2009	C
12345	Shankar	Comp. Sci.	32	CS-190	2	Spring	2009	A
12345	Shankar	History	32	CS-315	1	Spring	2010	A
12345	Shankar	Finance	32	CS-347	1	Fall	2009	A
19991	Brandt	Music	80	HIS-351	1	Spring	2010	B
23121	Chavez	Physics	110	FIN-201	1	Spring	2010	C+
44553	Peltier	Physics	56	PHY-101	1	Fall	2009	B-
45678	Levy	Physics	46	CS-101	1	Fall	2009	F
45678	Levy	Physics	46	CS-101	1	Spring	2010	B+
45678	Levy	Physics	46	CS-319	1	Spring	2010	B
54321	Williams	Comp. Sci.	54	CS-101	1	Fall	2009	A-
54321	Williams	Comp. Sci.	54	CS-190	2	Spring	2009	B+
55739	Sanchez	Music	38	MU-199	1	Spring	2010	A-
76543	Brown	Comp. Sci.	58	CS-101	1	Fall	2009	A
76543	Brown	Comp. Sci.	58	CS-319	2	Spring	2010	A
76653	Aoi	Elec. Eng.	60	EE-181	1	Spring	2009	C
98765	Bourikas	Elec. Eng.	98	CS-101	1	Fall	2009	C-
98765	Bourikas	Elec. Eng.	98	CS-315	1	Spring	2010	B
98988	Tanaka	Biology	120	BIO-101	1	Summer	2009	A
98988	Tanaka	Biology	120	BIO-301	1	Summer	2010	null

Figure 4.04

ID	name	dept_name	tot_cred	course_id	sec_id	semester	year	grade
00128	Zhang	Comp. Sci.	102	CS-101	1	Fall	2009	A
00128	Zhang	Comp. Sci.	102	CS-347	1	Fall	2009	A-
12345	Shankar	Comp. Sci.	32	CS-101	1	Fall	2009	C
12345	Shankar	Comp. Sci.	32	CS-190	2	Spring	2009	A
12345	Shankar	History	32	CS-315	1	Spring	2010	A
12345	Shankar	Finance	32	CS-347	1	Fall	2009	A
19991	Brandt	Music	80	HIS-351	1	Spring	2010	B
23121	Chavez	Physics	110	FIN-201	1	Spring	2010	C+
44553	Peltier	Physics	56	PHY-101	1	Fall	2009	B-
45678	Levy	Physics	46	CS-101	1	Fall	2009	F
45678	Levy	Physics	46	CS-101	1	Spring	2010	B+
45678	Levy	Physics	46	CS-319	1	Spring	2010	B
54321	Williams	Comp. Sci.	54	CS-101	1	Fall	2009	A-
54321	Williams	Comp. Sci.	54	CS-190	2	Spring	2009	B+
55739	Sanchez	Music	38	MU-199	1	Spring	2010	A-
70557	Snow	Physics	0	null	null	null	null	null
76543	Brown	Comp. Sci.	58	CS-101	1	Fall	2009	A
76543	Brown	Comp. Sci.	58	CS-319	2	Spring	2010	A
76653	Aoi	Elec. Eng.	60	EE-181	1	Spring	2009	C
98765	Bourikas	Elec. Eng.	98	CS-101	1	Fall	2009	C-
98765	Bourikas	Elec. Eng.	98	CS-315	1	Spring	2010	B
98988	Tanaka	Biology	120	BIO-101	1	Summer	2009	A
98988	Tanaka	Biology	120	BIO-301	1	Summer	2010	null

Figure 4.05

ID	course_id	sec_id	semester	year	grade	name	dept_name	tot_cred
00128	CS-101	1	Fall	2009	A	Zhang	Comp. Sci.	102
00128	CS-347	1	Fall	2009	A-	Zhang	Comp. Sci.	102
12345	CS-101	1	Fall	2009	C	Shankar	Comp. Sci.	32
12345	CS-190	2	Spring	2009	A	Shankar	Comp. Sci.	32
12345	CS-315	1	Spring	2010	A	Shankar	History	32
12345	CS-347	1	Fall	2009	A	Shankar	Finance	32
19991	HIS-351	1	Spring	2010	B	Brandt	Music	80
23121	FIN-201	1	Spring	2010	C+	Chavez	Physics	110
44553	PHY-101	1	Fall	2009	B-	Peltier	Physics	56
45678	CS-101	1	Fall	2009	F	Levy	Physics	46
45678	CS-101	1	Spring	2010	B+	Levy	Physics	46
45678	CS-319	1	Spring	2010	B	Levy	Physics	46
54321	CS-101	1	Fall	2009	A-	Williams	Comp. Sci.	54
54321	CS-190	2	Spring	2009	B+	Williams	Comp. Sci.	54
55739	MU-199	1	Spring	2010	A-	Sanchez	Music	38
70557	null	null	null	null	null	Snow	Physics	0
76543	CS-101	1	Fall	2009	A	Brown	Comp. Sci.	58
76543	CS-319	2	Spring	2010	A	Brown	Comp. Sci.	58
76653	EE-181	1	Spring	2009	C	Aoi	Elec. Eng.	60
98765	CS-101	1	Fall	2009	C-	Bourikas	Elec. Eng.	98
98765	CS-315	1	Spring	2010	B	Bourikas	Elec. Eng.	98
98988	BIO-101	1	Summer	2009	A	Tanaka	Biology	120
98988	BIO-301	1	Summer	2010	null	Tanaka	Biology	120

Figure 4.07

ID	name	dept_name	salary
10101	Srinivasan	Comp. Sci.	65000
12121	Wu	Finance	90000
15151	Mozart	Music	40000
22222	Einstein	Physics	95000
32343	El Said	History	60000
33456	Gold	Physics	87000
45565	Katz	Comp. Sci.	75000
58583	Califieri	History	62000
76543	Singh	Finance	80000
76766	Crick	Biology	72000
83821	Brandt	Comp. Sci.	92000
98345	Kim	Elec. Eng.	80000
69987	White	*null*	*null*

instructor

dept_name	building	budget
Biology	Watson	90000
Comp. Sci.	Taylor	100000
Elec. Eng.	Taylor	85000
Finance	Painter	120000
History	Painter	50000
Music	Packard	80000
Physics	Watson	70000
null	Taylor	*null*

department

Figure 4.06

Join types
inner join
left outer join
right outer join
full outer join

Join conditions
natural
on $<$predicate$>$
using $(A_1, A_2, \ldots, A_n)$

Figure 4.03

CS425 – Fall 2013
Boris Glavic
Chapter 6: Advanced SQL

Chapter 6: Advanced SQL

- Accessing SQL From a Programming Language
  - Dynamic SQL
    - JDBC and ODBC
  - Embedded SQL
- Functions and Procedural Constructs
- Triggers

Textbook: Chapter 5

Accessing SQL From a Programming Language

JDBC and ODBC

- API (application-program interface) for a program to interact with a database server

- Application makes calls to
  - Connect with the database server
  - Send SQL commands to the database server
  - Fetch tuples of result one-by-one into program variables

- ODBC (Open Database Connectivity) works with C, C++, C#, and Visual Basic
  - Other API's such as ADO.NET sit on top of ODBC

- JDBC (Java Database Connectivity) works with Java

Native APIs

- Most DBMS also define DBMS specific APIs

- Oracle: OCI
- Postgres: libpg

…

JDBC

- **JDBC** is a Java API for communicating with database systems supporting SQL.

- JDBC supports a variety of features for querying and updating data, and for retrieving query results.

- JDBC also supports metadata retrieval, such as querying about relations present in the database and the names and types of relation attributes.

- Model for communicating with the database:

  - Open a connection

  - Create a "statement" object

  - Execute queries using the Statement object to send queries and fetch results

  - Exception mechanism to handle errors

JDBC Code

```
public static void JDBCexample(String dbid, String userid, String passwd)
{
    try {
        Class.forName ("oracle.jdbc.driver.OracleDriver"); // load driver
        Connection conn = DriverManager.getConnection(  // connect to server
            "jdbc:oracle:thin:@db.yale.edu:2000:univdb", userid, passwd);
        Statement stmt = conn.createStatement(); // create Statement object
            … Do Actual Work ….
        stmt.close(); // close Statement and release resources
        conn.close(); // close Connection and release resources
    }
    catch (SQLException sqle) {
        System.out.println("SQLException : " + sqle); // handle exceptions
    }
}
```

JDBC Code (Cont.)

- Update to database
```
try {
    stmt.executeUpdate(
        "insert into instructor values('77987', 'Kim', 'Physics',
98000)");
} catch (SQLException sqle)
{
    System.out.println("Could not insert tuple. " + sqle);
}
```

- Execute query and fetch and print results
```
        ResultSet rset = stmt.executeQuery(
                        "select dept_name, avg (salary)
                         from instructor
                         group by dept_name");
    while (rset.next()) {
        System.out.println(rset.getString("dept_name") + " " +
                                rset.getFloat(2));
    }
```

JDBC Code Details

- Result stores the current row position in the result
    - Pointing before the first row after executing the statement
    - **.next()** moves to the next tuple
        - Returns false if no more tuples

- Getting result fields:

    - **rs.getString("dept_name") and rs.getString(1) equivalent if dept_name is the first attribute in select result.**

- Dealing with Null values

    - **int a = rs.getInt("a");**

        if (rs.wasNull()) Systems.out.println("Got null value");

Prepared Statement

■ PreparedStatement pStmt = conn.prepareStatement(
                              "insert into instructor values(?,?,?,?)");
pStmt.setString(1, "88877");   pStmt.setString(2, "Perry");
pStmt.setString(3, "Finance");   pStmt.setInt(4, 125000);
pStmt.executeUpdate();
pStmt.setString(1, "88878");
pStmt.executeUpdate();

■ For queries, use pStmt.executeQuery(), which returns a ResultSet

■ WARNING: always use prepared statements when taking an input from the user and adding it to a query

  ● **NEVER create a query by concatenating strings which you get as inputs**

  ● "insert into instructor values(' " + ID + " ' , ' " + name + " ' , " +
                                 " ' " + dept name + " ' , " ' balance +
    ")"

  ● What if name is "D' Souza"?

SQL Injection

■ Suppose query is constructed using

● "select * from instructor where name = ' " + name + "' "

■ Suppose the user, instead of entering a name, enters:

● X' or ' Y' = ' Y

■ then the resulting statement becomes:

● "select * from instructor where name = ' " + "X' or ' Y' = ' Y" + "' "

● which is:

▸ select * from instructor where name = ' X' or ' Y' = ' Y'

● User could have even used

▸ X' ; update instructor set salary = salary + 10000; --

■ Prepared statement internally uses:
"select * from instructor where name = ' X\' or \' Y\' = \' Y'

● **Always use prepared statements, with user inputs as parameters**

Metadata Features

- ResultSet metadata

- E.g., after executing query to get a ResultSet rs:

  - ResultSetMetaData rsmd = rs.getMetaData();

    for(int i = 1; i <= rsmd.getColumnCount(); i++) {

      System.out.println(rsmd.getColumnName(i));

      System.out.println(rsmd.getColumnTypeName(i));

    }

- How is this useful?

Metadata (Cont)

- Database metadata

- DatabaseMetaData dbmd = conn.getMetaData();

  ResultSet rs = dbmd.getColumns(null, "univdb", "department", "%");

  // Arguments to getColumns: Catalog, Schema-pattern, Table-pattern,

  // and Column-Pattern

  // Returns: One row for each column; row has a number of attributes

  // such as COLUMN_NAME, TYPE_NAME

  while(rs.next()) {

      System.out.println(rs.getString("COLUMN_NAME"),

                             rs.getString("TYPE_NAME");

  }

- And where is this useful?

Transaction Control in JDBC

■ By default, each SQL statement is treated as a separate transaction that is committed automatically

  ● bad idea for transactions with multiple updates

■ Can turn off automatic commit on a connection

  ● conn.setAutoCommit(false);

■ Transactions must then be committed or rolled back explicitly

  ● conn.commit();   or

  ● conn.rollback();

■ conn.setAutoCommit(true) turns on automatic commit.

Other JDBC Features

■ Calling functions and procedures

- CallableStatement cStmt1 = conn.prepareCall("{? = call some function(?)}");

- CallableStatement cStmt2 = conn.prepareCall("{call some procedure(?,?)}");

■ Handling large object types

- getBlob() and getClob() that are similar to the getString() method, but return objects of type Blob and Clob, respectively

- get data from these objects by getBytes()

- associate an open stream with Java Blob or Clob object to update large objects

  ‣ blob.setBlob(int parameterIndex, InputStream inputStream).

SQLJ

- JDBC is overly dynamic, errors cannot be caught by compiler
- SQLJ: embedded SQL in Java
  - #sql iterator deptInfoIter (String dept name, int avgSal);

    deptInfoIter iter = null;

    #sql iter = { select dept_name, avg(salary) from instructor

                  group by dept name };

    while (iter.next()) {

        String deptName = iter.dept_name();

        int avgSal = iter.avgSal();

        System.out.println(deptName + " " + avgSal);

    }

    iter.close();

ODBC

- Open DataBase Connectivity(ODBC) standard

  - standard for application program to communicate with a database server.

  - application program interface (API) to

    - open a connection with a database,

    - send queries and updates,

    - get back results.

- Applications such as GUI, spreadsheets, etc. can use ODBC

- Was defined originally for Basic and C, versions available for many languages.

ODBC (Cont.)

- Each database system supporting ODBC provides a "driver" library that must be linked with the client program.

- When client program makes an ODBC API call, the code in the library communicates with the server to carry out the requested action, and fetch results.

- ODBC program first allocates an SQL environment, then a database connection handle.

- Opens database connection using SQLConnect(). Parameters for SQLConnect:

  - connection handle,

  - the server to which to connect

  - the user identifier,

  - password

- Must also specify types of arguments:

  - SQL_NTS denotes previous argument is a null-terminated string.

ODBC Code

- **int ODBCexample()**

  {

      RETCODE error;

      HENV   env;   /* environment */

      HDBC   conn; /* database connection */

      SQLAllocEnv(&env);

      SQLAllocConnect(env, &conn);

      SQLConnect(conn, "db.yale.edu", SQL_NTS, "avi", SQL_NTS,
        "avipasswd", SQL_NTS);

      { …. Do actual work … }

      SQLDisconnect(conn);

      SQLFreeConnect(conn);

      SQLFreeEnv(env);

  }

ODBC Code (Cont.)

- Program sends SQL commands to database by using SQLExecDirect
- Result tuples are fetched using SQLFetch()
- SQLBindCol() binds C language variables to attributes of the query result
  - When a tuple is fetched, its attribute values are automatically stored in corresponding C variables.
  - Arguments to SQLBindCol()
    - ODBC stmt variable, attribute position in query result
    - The type conversion from SQL to C.
    - The address of the variable.
    - For variable-length types like character arrays,
      - The maximum length of the variable
      - Location to store actual length when a tuple is fetched.
      - Note: A negative value returned for the length field indicates null value
- Good programming requires checking results of every function call for errors; we have omitted most checks for brevity.

ODBC Code (Cont.)

■ Main body of program

```
 char deptname[80];
float salary;
int lenOut1, lenOut2;
HSTMT stmt;
char * sqlquery = "select dept_name, sum (salary)
                          from instructor
                          group by dept_name";
SQLAllocStmt(conn, &stmt);
error = SQLExecDirect(stmt, sqlquery, SQL_NTS);
if (error == SQL SUCCESS) {
       SQLBindCol(stmt, 1, SQL_C_CHAR, deptname , 80, &lenOut1);
       SQLBindCol(stmt, 2, SQL_C_FLOAT, &salary, 0 , &lenOut2);
       while (SQLFetch(stmt) == SQL_SUCCESS) {
              printf (" %s %g\n", deptname, salary);
       }
}
SQLFreeStmt(stmt, SQL_DROP);
```

ODBC Prepared Statements

- **Prepared Statement**
  - SQL statement prepared: compiled at the database
  - Can have placeholders: E.g. insert into account values(?,?,?)
  - Repeatedly executed with actual values for the placeholders
- To prepare a statement
  SQLPrepare(stmt, <SQL String>);
- To bind parameters
  SQLBindParameter(stmt, <parameter#>,
        … type information and value omitted for simplicity..)
- To execute the statement
  retcode = SQLExecute(stmt);
- To avoid SQL injection security risk, do not create SQL strings directly using user input; instead use prepared statements to bind user inputs

More ODBC Features

- **Metadata features**
  - finding all the relations in the database and
  - finding the names and types of columns of a query result or a relation in the database.

- By default, each SQL statement is treated as a separate transaction that is committed automatically.

  - Can turn off automatic commit on a connection
    - SQLSetConnectOption(conn, SQL_AUTOCOMMIT, 0)}

  - Transactions must then be committed or rolled back explicitly by
    - SQLTransact(conn, SQL_COMMIT) or
    - SQLTransact(conn, SQL_ROLLBACK)

ODBC Conformance Levels

- Conformance levels specify subsets of the functionality defined by the standard.

  - Core

  - Level 1 requires support for metadata querying

  - Level 2 requires ability to send and retrieve arrays of parameter values and more detailed catalog information.

- SQL Call Level Interface (CLI) standard similar to ODBC interface, but with some minor differences.

ADO.NET

- API designed for Visual Basic .NET and C#, providing database access facilities similar to JDBC/ODBC

  - Partial example of ADO.NET code in C#
    ```
    using System, System.Data, System.Data.SqlClient;
    SqlConnection conn = new SqlConnection(
                "Data Source=<IPaddr>, Initial Catalog=<Catalog>");
    conn.Open();
    SqlCommand cmd = new SqlCommand("select * from students",
                                        conn);
    SqlDataReader rdr = cmd.ExecuteReader();
    while(rdr.Read()) {
        Console.WriteLine(rdr[0], rdr[1]); /* Prints result attributes 1 & 2 */
    }
    rdr.Close(); conn.Close();
    ```

- Can also access non-relational data sources such as

  - OLE-DB, XML data, Entity framework

Dynamic vs. Embedded SQL

Dynamic SQL

code

Compiler

binary

Library

DBMS

Embedded SQL

Code with embeded SQL

Preprocessor → code

Compiler

binary

Library

DBMS

Embedded SQL

- The SQL standard defines embeddings of SQL in a variety of programming languages such as C, Java, and Cobol.

- A language to which SQL queries are embedded is referred to as a **host language**, and the SQL structures permitted in the host language comprise *embedded* SQL.

- The basic form of these languages follows that of the System R embedding of SQL into PL/I.

- **EXEC SQL** statement is used to identify embedded SQL request to the preprocessor

    EXEC SQL <embedded SQL statement > END_EXEC

    Note: this varies by language (for example, the Java embedding uses   # SQL { …. };)

Example Query

- From within a host language, find the ID and name of students who have completed more than the number of credits stored in variable credit_amount.

- Specify the query in SQL and declare a *cursor* for it

  EXEC SQL

  declare *c* **cursor for**
  select *ID, name*
  from *student*
  where tot_cred *> :credit_amount*

  END_EXEC

Embedded SQL (Cont.)

■ The **open** statement causes the query to be evaluated

EXEC SQL **open** *c* END_EXEC

■ The **fetch** statement causes the values of one tuple in the query result to be placed on host language variables.

EXEC SQL **fetch** *c* **into** :*si, :sn* END_EXEC

Repeated calls to **fetch** get successive tuples in the query result

■ A variable called SQLSTATE in the SQL communication area (SQLCA) gets set to '02000' to indicate no more data is available

■ The **close** statement causes the database system to delete the temporary relation that holds the result of the query.

EXEC SQL **close** *c* END_EXEC

Note: above details vary with language. For example, the Java embedding defines Java iterators to step through result tuples.

Updates Through Cursors

■ Can update tuples fetched by cursor by declaring that the cursor is for update

> **declare** *c* **cursor for**
>   **select** *
>   **from** *instructor*
>   **where** *dept_name* = 'Music'
> **for update**

■ To update tuple at the current location of cursor *c*

> **update** *instructor*
> **set** *salary = salary* + 100
> **where current of** *c*

Procedural Constructs in SQL

Procedural Extensions and Stored Procedures

■ SQL provides a **module** language

- Permits definition of procedures in SQL, with if-then-else statements, for and while loops, etc.

■ Stored Procedures

- Can store procedures in the database

- then execute them using the **call** statement

- permit external applications to operate on the database without knowing about internal details

■ Object-oriented aspects of these features are covered in Chapter 22 (Object Based Databases) in the textbook

Why have procedural extensions?

- Shipping data between a database server and application program (e.g., through network connection) is costly

- Converting data from the database internal format into a format understood by the application programming language is costly

- Example:
  - Use Java to retrieve all users and their friend-relationships from a friends relation representing a world-wide social network with 10,000,000 users

  - Compute the transitive closure
    - All pairs of users connects through a path of friend relationships. E.g., (Peter, Magret) if Peter is a friend of Walter who is a friend of Magret

  - Return pairs of users from Chicago – say 4000 pairs

  - 1) cannot be expressed (efficiently) as SQL query, 2) result is small
    - -> save by executing this on the DB server

Functions and Procedures

- SQL:1999 supports functions and procedures
  - Functions/procedures can be written in SQL itself, or in an external programming language.
  - Functions are particularly useful with specialized data types such as images and geometric objects.
    - Example: functions to check if polygons overlap, or to compare images for similarity.
  - Some database systems support **table-valued functions**, which can return a relation as a result.
- SQL:1999 also supports a rich set of imperative constructs, including
  - Loops, if-then-else, assignment
- Many databases have proprietary procedural extensions to SQL that differ from SQL:1999.

SQL Functions

■ Define a function that, given the name of a department, returns the count of the number of instructors in that department.

> **create function** *dept_count* (*dept_name* **varchar**(20))
> **returns integer**
> **begin**
>   **declare** *d_count* **integer;**
>   **select count** (***) **into** *d_count*
>   **from** *instructor*
>   **where** *instructor.dept_name = dept_name;*
>   **return** *d_count;*
> **end**

■ Find the department name and budget of all departments with more that 12 instructors.

> **select** *dept_name, budget*
> **from** *department*
> **where** *dept_*count (*dept_name*) > 1

Table Functions

- SQL:2003 added functions that return a relation as a result

- Example: Return all accounts owned by a given customer

create function *instructors_of* (*dept_name* **char**(20)

      returns table (*ID* **varchar**(5),
                          name **varchar**(20),
                          dept_name **varchar**(20),
                          salary **numeric**(8,2))

return table
    (**select** *ID, name, dept_name, salary*
     from *instructor*
     where *instructor.dept_name = instructors_of.dept_name*)

- Usage

    select *
    from table (*instructors_of* ('Music'))

SQL Procedures

■ The *dept_count* function could instead be written as procedure:

create procedure *dept_count_proc* (**in** *dept_name* **varchar**(20),
out *d_count* **integer)**

begin

  select count(***) **into** *d_count*
  from *instructor*
  where *instructor.dept_name = dept_count_proc.dept_name*

end

■ Procedures can be invoked either from an SQL procedure or from embedded SQL, using the **call** statement.

      declare *d_count* **integer**;
      call *dept_count_proc*('Physics' , *d_count*);

Procedures and functions can be invoked also from dynamic SQL

■ SQL:1999 allows more than one function/procedure of the same name (called name **overloading**), as long as the number of arguments differ, or at least the types of the arguments differ

Procedural Constructs

- Warning: most database systems implement their own variant of the standard syntax below

  - read your system manual to see what works on your system

- Compound statement: **begin ... end**,

  - May contain multiple SQL statements between **begin** and **end**.

  - Local variables can be declared within a compound statements

- **While** and **repeat** statements :

> **declare** n **integer default** 0;
> **while** $n < 10$ **do**
>   **set** $n = n + 1$
> **end while**
>
> **repeat**
>   **set** $n = n - 1$
> **until** $n = 0$
> **end repeat**

Procedural Constructs (Cont.)

- **For** loop
  - Permits iteration over all results of a query
  - Example:

      declare *n* **integer default** 0;
      for *r* **as**
          select *budget* **from** *department*
           where *dept_name* = 'Music'
       do
          set *n* = *n* - r.*budget*
       end for

Procedural Constructs (cont.)

- Conditional statements (**if-then-else**)
  SQL:1999 also supports a **case** statement similar to C case statement

- Example procedure: registers student after ensuring classroom capacity is not exceeded

  - Returns 0 on success and -1 if capacity is exceeded

  - See book for details

- Signaling of exception conditions, and declaring handlers for exceptions

  > **declare** *out_of_classroom_seats* **condition**
  > **declare exit handler for** *out_of_classroom_seats*
  > **begin**
  > …
  > .. **signal** *out_of_classroom_seats*
  > **end**

  - The handler here is **exit** -- causes enclosing **begin..end** to be exited

  - Other actions possible on exception

External Language Functions/Procedures

- SQL:1999 permits the use of functions and procedures written in other languages such as C or C++

- Declaring external language procedures and functions

create procedure dept_count_proc(**in** *dept_name* **varchar**(20),
                                    out count **integer**)

language C
external name ' /usr/avi/bin/dept_count_proc'

create function dept_count(*dept_name* **varchar**(20))
returns integer
language C
external name '/usr/avi/bin/dept_count'

External Language Routines (Cont.)

- **Benefits of external language functions/procedures:**
  - more efficient for many operations, and more expressive power.

- **Drawbacks**
  - Code to implement function may need to be loaded into database system and executed in the database system's address space.
    - ‣ risk of accidental corruption of database structures
    - ‣ security risk, allowing users access to unauthorized data
  - There are alternatives, which give good security at the cost of potentially worse performance.
  - Direct execution in the database system's space is used when efficiency is more important than security.

Security with External Language Routines

- To deal with security problems
  - Use **sandbox** techniques
    - ▸ E.g., use a safe language like Java, which cannot be used to   access/damage other parts of the database code.
  - Or, run external language functions/procedures in a separate process, with no access to the database process' memory.
    - ▸ Parameters and results communicated via inter-process communication
- Both have performance overheads
- Many database systems support both above approaches as well as direct executing in database system address space.

Triggers

Triggers

- A **trigger** is a statement that is executed automatically by the system as a side effect of a modification to the database.

- To design a trigger mechanism, we must:
  - Specify the conditions under which the trigger is to be executed.
  - Specify the actions to be taken when the trigger executes.

- Triggers introduced to SQL standard in SQL:1999, but supported even earlier using non-standard syntax by most databases.
  - Syntax illustrated here may not work exactly on your database system; check the system manuals

Trigger Example

■ E.g. *time_slot_id* is not a primary key of *timeslot,* so we cannot create a foreign key constraint from *section* to *timeslot.*

■ Alternative: use triggers on *section* and *timeslot* to enforce integrity constraints

create trigger *timeslot_check1* **after insert on** *section*
referencing new row as *nrow*
for each row
when (*nrow.time_slot_id* **not in** (
        select *time_slot_id*
        from *time_slot*)) /* *time_slot_id* not present in *time_slot* */
begin
   rollback
end;

Trigger Example Cont.

create trigger *timeslot_check2* **after delete on** *timeslot*
    referencing old row as *orow*
    for each row
    when (*orow.time_slot_id* **not in** (
            select *time_slot_id*
            from *time_slot*)
            /* last tuple for *time slot id* deleted from *time slot* */
        and *orow.time_slot_id* **in** (
            select *time_slot_id*
            from *section*)) /* and *time_slot_id* still referenced from *section*/
    begin
      rollback
    end;

Triggering Events and Actions in SQL

- Triggering event can be **insert**, **delete** or **update**

- Triggers on update can be restricted to specific attributes
  - **E.g., after update of** *takes* **on** *grade*

- Values of attributes before and after an update can be referenced
  - **referencing old row as** **:** for deletes and updates
  - **referencing new row as** **:** for inserts and updates

- Triggers can be activated before an event, which can serve as extra constraints. E.g. convert blank grades to null.

  create trigger *setnull_trigger* **before update of** *takes*
  referencing new row as *nrow*
  for each row
  when (*nrow.grade* = ' ')
   begin atomic
      set *nrow.grade* = **null;**
   end;

Trigger to Maintain credits_earned value

- **create trigger** *credits_earned* **after update of** *takes* **on** (*grade*)
  referencing new row as *nrow*
  referencing old row as *orow*
  for each row
  when *nrow.grade* <> 'F' **and** *nrow.grade* **is not null**
    and (*orow.grade* = 'F' **or** *orow.grade* **is null**)
  begin atomic
     update *student*
     set *tot_cred*= *tot_cred* +
         (**select** *credits*
          from *course*
          where *course.course_id*= *nrow.course_id*)
     where *student.id* = *nrow.id*;
  end;

Statement Level Triggers

■ Instead of executing a separate action for each affected row, a single action can be executed for all rows affected by a transaction

- Use **for each statement** instead of **for each row**

- Use **referencing old table** or **referencing new table** to refer to temporary tables (called *transition tables*) containing the affected rows

- Can be more efficient when dealing with SQL statements that update a large number of rows

When Not To Use Triggers

- Triggers were used earlier for tasks such as

  - maintaining summary data (e.g., total salary of each department)

  - Replicating databases by recording changes to special relations (called **change** or **delta** relations) and having a separate process that applies the changes over to a replica

- There are better ways of doing these now:

  - Databases today provide built in materialized view facilities to maintain summary data

  - Databases provide built-in support for replication

- Encapsulation facilities can be used instead of triggers in many cases

  - Define methods to update fields

  - Carry out actions as part of the update methods instead of through a trigger

When Not To Use Triggers

- Risk of unintended execution of triggers, for example, when

  - loading data from a backup copy

  - replicating updates at a remote site

  - Trigger execution can be disabled before such actions.

- Other risks with triggers:

  - Error leading to failure of critical transactions that set off the trigger

  - Cascading execution

Recursive Queries

Recursion in SQL

- SQL:1999 permits recursive view definition

- Example: find which courses are a prerequisite, whether directly or indirectly, for a specific course

  with recursive *rec_prereq*(*course_id*, *prereq_id*) **as** (
      select *course_id*, *prereq_id*
      from *prereq*
    union
      select *rec_prereq.course_id*, *prereq.prereq_id*,
      from *rec_rereq*, *prereq*
      where *rec_prereq.prereq_id = prereq.course_id*
)
  select $*$
  from *rec_prereq*;

  This example view, *rec_prereq,* is called the *transitive closure* of the *prereq* relation

The Power of Recursion

■ Recursive views make it possible to write queries, such as transitive closure queries, that cannot be written without recursion or iteration.

- ● Intuition: Without recursion, a non-recursive non-iterative program can perform only a fixed number of joins of *prereq* with itself

  - ▸ This can give only a fixed number of levels of managers

  - ▸ Given a fixed non-recursive query, we can construct a database with a greater number of levels of prerequisites on which the query will not work

  - ▸ Alternative: write a procedure to iterate as many times as required

    – See procedure *findAllPrereqs* in book

The Power of Recursion

- Computing transitive closure using iteration, adding successive tuples to *rec_prereq*

  - The next slide shows a *prereq* relation

  - Each step of the iterative process constructs an extended version of *rec_prereq* from its recursive definition.

  - The final result is called the *fixed point* of the recursive view definition.

- Recursive views are required to be **monotonic**. That is,

  - if we add tuples to *prereq* the view *rec_prereq* contains all of the tuples it contained before, plus possibly more

course_id	prereq_id
BIO-301	BIO-101
BIO-399	BIO-101
CS-190	CS-101
CS-315	CS-101
CS-319	CS-101
CS-347	CS-101
EE-181	PHY-101

Iteration Number	Tuples in cl
0	
1	(CS-301)
2	(CS-301), (CS-201)
3	(CS-301), (CS-201)
4	(CS-301), (CS-201), (CS-101)
5	(CS-301), (CS-201), (CS-101)

Another Recursion Example

- Given relation

    manager(*employee_name, manager_name*)

- Find all employee-manager pairs, where the employee reports to the manager directly or indirectly (that is manager's manager, manager's manager's manager, etc.)

    with recursive *empl* (*employee_name, manager_name*) **as** (
          select *employee_name, manager_name*
          from   *manager*
       union
          select manager.*employee_name*, empl.*manager_name*
          from   *manager*, *empl*
          where *manager.manager_name = empl.emploe_name*)
    select *
    from   *empl*

    This example view, *empl,* is the *transitive closure* of the *manager* relation

Recap

- Programming Language Interfaces for Databases
  - Dynamic SQL (e.g., JDBC, ODBC)
  - Embedded SQL
  - SQL Injection
- Procedural Extensions of SQL
  - Functions and Procedures
- External Functions/Procedures
  - Written in programming language (e.g., C)
- Triggers
  - Events (insert, …)
  - Conditions (WHEN)
  - per statement / per row
  - Accessing old/new table/row versions
- Recursive Queries

End of Chapter

Outline

- Introduction
- Relational Data Model
- Formal Relational Languages (relational algebra)
- SQL - Advanced
- **Database Design – ER model**
- Transaction Processing, Recovery, and Concurrency Control
- Storage and File Structures
- Indexing and Hashing
- Query Processing and Optimization

CS425 – Fall 2013
Boris Glavic
Chapter 7: Entity-Relationship Model

Partially taken from

Klaus R. Dittrich

modified from:

Database System Concepts, 6th Ed.

Chapter 7: Entity-Relationship Model

- Design Process
- Modeling
- Constraints
- E-R Diagram
- Design Issues
- Weak Entity Sets
- Extended E-R Features
- Design of the Bank Database
- Reduction to Relation Schemas
- Database Design
- UML

Database Design

World

???

Relational DB schema
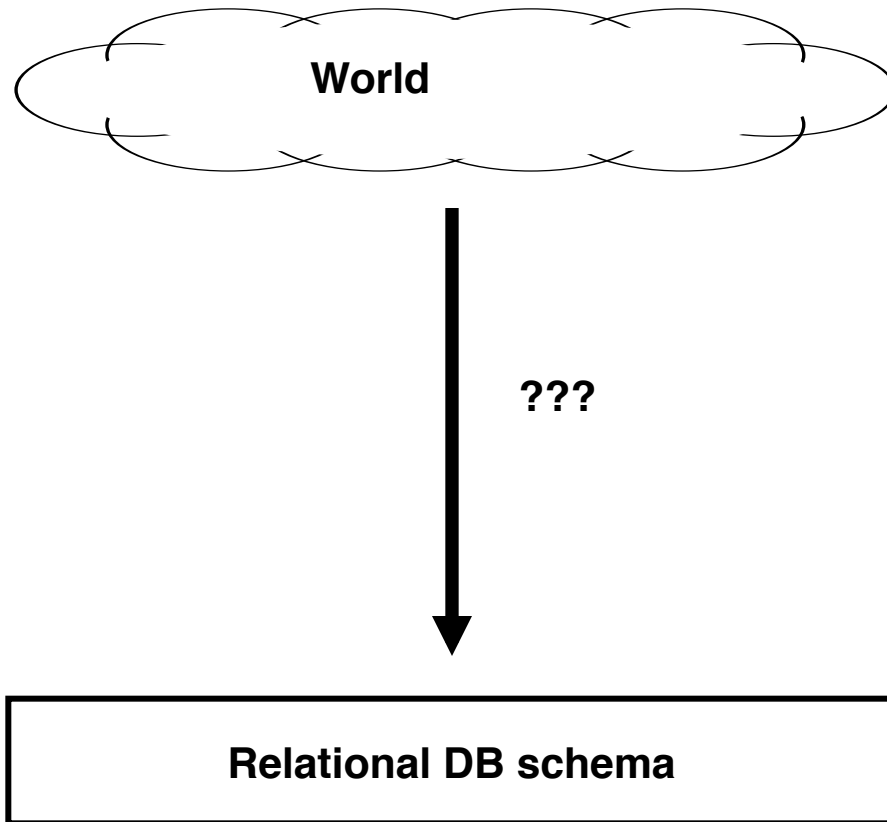
Database Design

First: need to develop a "mind"-model based on a requirement analysis

World

Requirement Analysis

"Mind" Model

English (e.g.)

???

Relational DB schema
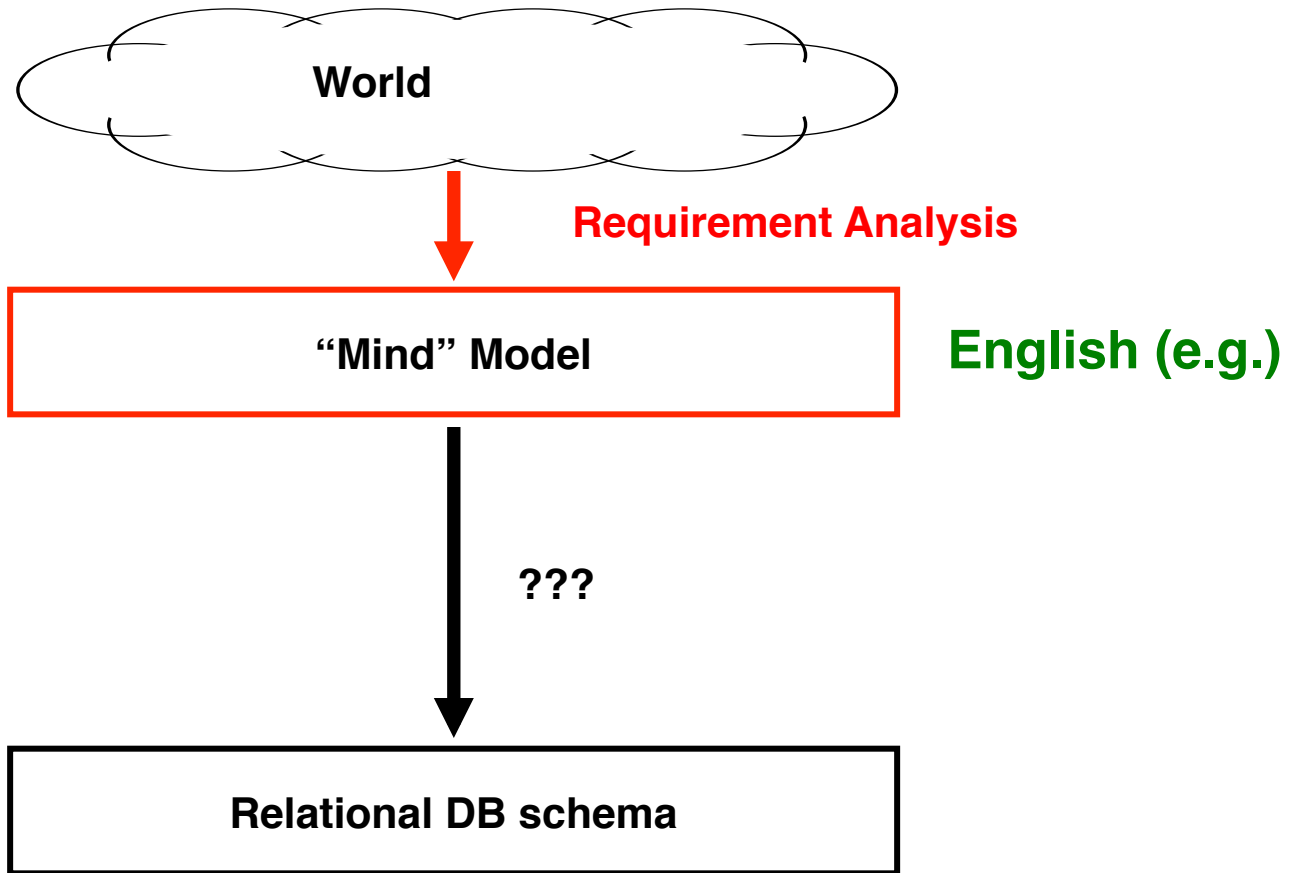
Requirement Analysis Example Zoo

- The zoo stores information about animals, cages, and zoo keepers.

- Animals are of a certain species and have a name. For each animal we want to record its weight and age.

- Each cage is located in a section of the zoo. Cages can house animals, but there may be cages that are currently empty. Cages have a size in square meter.

- Zoo keepers are identified by their social security number. We store a first name, last name, and for each zoo keeper. Zoo keepers are assigned to cages they have to take care of (clean, …). Each cage that is not empty has a zoo keeper assigned to it. A zoo keeper can take care of several cages. Each zoo keeper takes care of at least one cage.

■ Let's do it!

Database Design

- Second: Formalize this model by developing a conceptual model

World

Requirement Analysis

"Mind" Model — **English (e.g.)**

Conceptual modeling

Conceptual Model — **ER model**

???

Relational DB schema
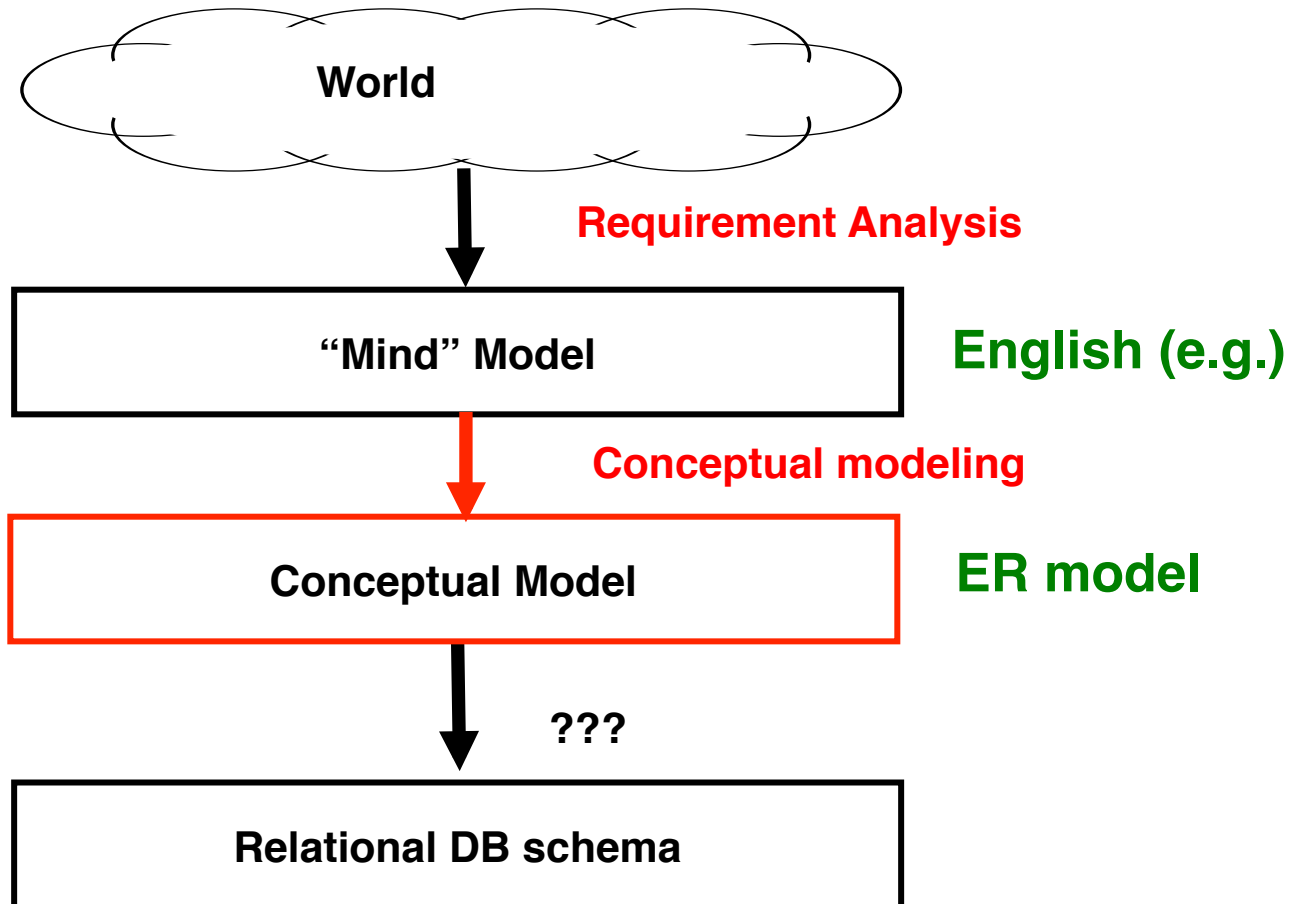
Database Design

■ Second: Formalize this model by developing a conceptual model

Modeling – ER model

- A *database* can be modeled as:
  - a collection of entities,
  - relationship among entities.
- An **entity** is an object that exists and is distinguishable from other objects.
  - Example: specific person, company, event, plant
- Entities have **attributes**
  - Example: people have *names* and *addresses*
- An **entity set** is a set of entities of the same type that share the same properties.
  - Example: set of all persons, companies, trees, holidays

Entity Sets *instructor* and *student*

instructor_ID instructor_name

76766	Crick
45565	Katz
10101	Srinivasan
98345	Kim
76543	Singh
22222	Einstein

instructor

student-ID student_name

98988	Tanaka
12345	Shankar
00128	Zhang
76543	Brown
76653	Aoi
23121	Chavez
44553	Peltier

student

Relationship Sets

■ A **relationship** is an association among several entities

Example:

44553 (Peltier)	_advisor_	22222 (Einstein)
student entity	relationship set	_instructor_ entity

■ A **relationship set** is a mathematical relation among $n \geq 2$ entities, each taken from entity sets

$$\{(e_1, e_2, \ldots e_n) \mid e_1 \in E_1, e_2 \in E_2, \ldots, e_n \in E_n\}$$

where $(e_1, e_2, \ldots, e_n)$ is a relationship

● Example:

$(44553, 22222) \in$ _advisor_

Relationship Set *advisor*

	instructor		student
76766	Crick	98988	Tanaka
45565	Katz	12345	Shankar
10101	Srinivasan	00128	Zhang
98345	Kim	76543	Brown
76543	Singh	76653	Aoi
22222	Einstein	23121	Chavez
		44553	Peltier

Relationship Sets (Cont.)

- An **attribute** can also be property of a relationship set.

- For instance, the *advisor* relationship set between entity sets *instructor* and *student* may have the attribute *date* which tracks when the student started being associated with the advisor

instructor
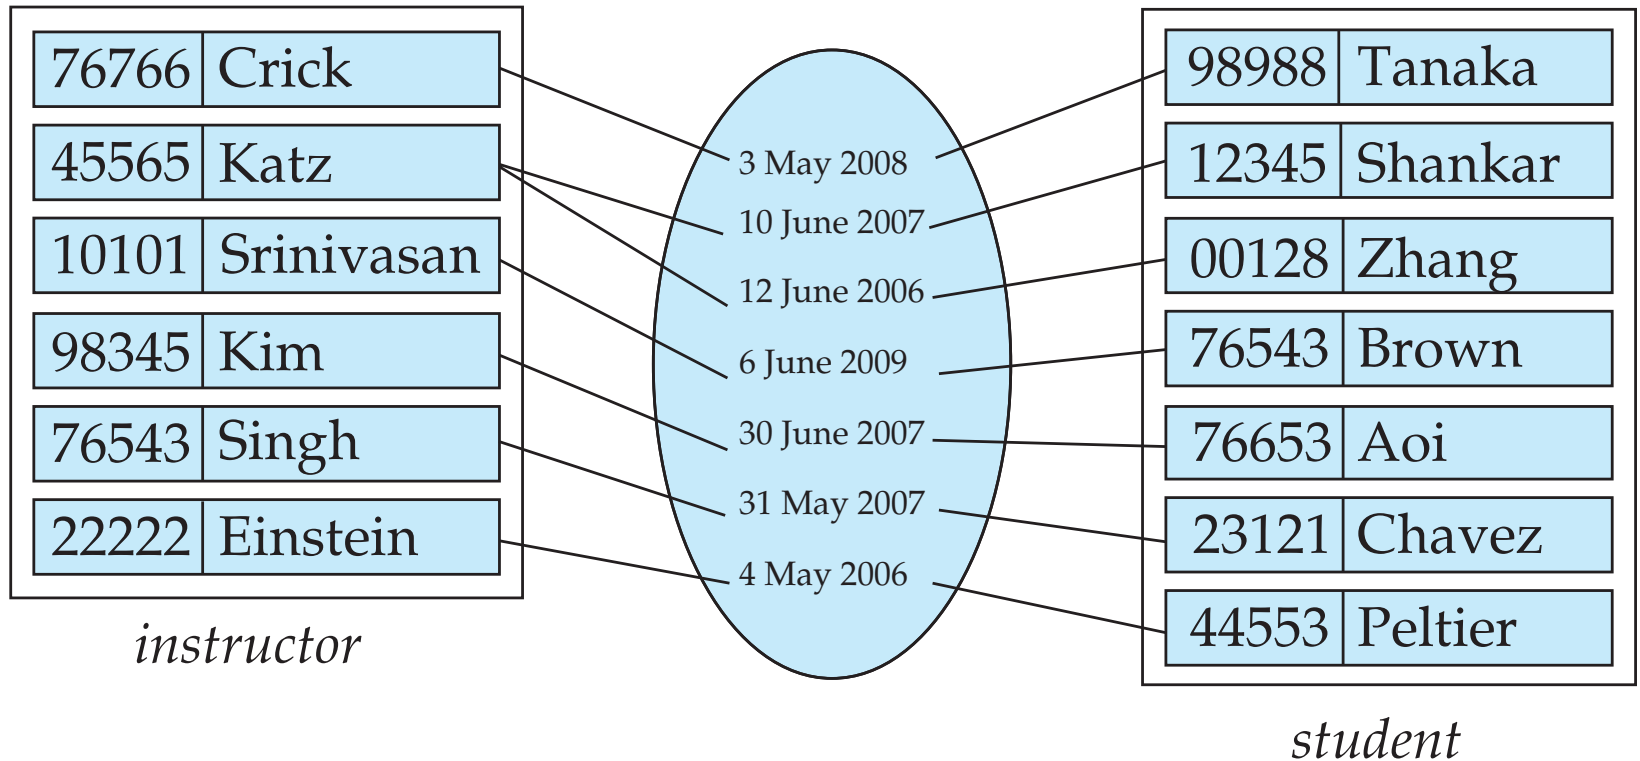
student

Degree of a Relationship Set

- **binary relationship**

  - involve two entity sets (or degree two).

- Relationships between more than two entity sets are rare. Most relationships are binary. (More on this later.)

  - Example: *students* work on research *projects* under the guidance of an *instructor*.

  - relationship *proj_guide* is a ternary relationship between *instructor, student,* and *project*

Attributes

- An entity is represented by a set of attributes, that are descriptive properties possessed by all members of an entity set.

    - Example:

        instructor = (ID, name, street, city, salary)
        course= (course_id, title, credits)

- **Domain** – the set of permitted values for each attribute

- Attribute types:

    - **Simple** and **composite** attributes.

    - **Single-valued** and **multivalued** attributes
        - Example: multivalued attribute: *phone_numbers*

    - **Derived** attributes
        - Can be computed from other attributes
        - Example: age, given date_of_birth

Composite Attributes

composite
attributes

component
attributes

```
              name                                    address
             /  |  \                              /    |    |     \
            /   |   \                            /     |    |      \
   first_name  middle_initial  last_name     street  city  state  postal_code
                                             /  |  \
                                            /   |   \
                            street_number  street_name  apartment_number
```

Mapping Cardinality Constraints

- Express the number of entities to which another entity can be associated via a relationship set.

- For a binary relationship set the mapping cardinality must be one of the following types:
  - One to one (1-1)
  - One to many (1-N)
  - Many to one (N-1)
  - Many to many (N-M)

(a) One to one

(b) One to many

Note: Some elements in *A* and *B* may not be mapped to any elements in the other set

Mapping Cardinalities Example

(a) One to one

(b) One to many

Note: Some elements in *A* and *B* may not be mapped to any elements in the other set

Mapping Cardinalities

(a)

Many to one

(b)

Many to many

Note: Some elements in A and B may not be mapped to any elements in the other set

Mapping Cardinalities Example

(a) Many to one

(b) Many to many

Note: Some elements in A and B may not be mapped to any elements in the other set

Mapping Cardinality Constraints Cont.

- What if we allow some elements to not be mapped to another element?

  - E.g., 0:1 – 1

- For a binary relationship set the mapping cardinality must be one of the following types:

- **1-1**
  - 1-1
  - 0:1-1
  - 1-0:1
  - 0:1-0:1

- **1-N**
  - 0:1-N
  - 0:1-0:N
  - 1-N
  - 1-0:N

- **N-1**
  - N-1
  - N-0:1
  - 0:N-1
  - 0:N-0:1

- **N-M**
  - N-M
  - N-0:M
  - 0:N-M
  - 0:N-0:M

Mapping Cardinality Constraints Cont.

- Typical Notation
  - (0:1) – (1:N)

Keys

- A **super key** of an entity set is a set of one or more attributes whose values uniquely determine each entity.

- A **candidate key** of an entity set is a minimal super key
  - *ID* is candidate key of *instructor*
  - *course_id* is candidate key of *course*

- Although several candidate keys may exist, one of the candidate keys is selected to be the **primary key**.

- **Note: Basically the same as for relational model**

Keys for Relationship Sets

■ The combination of primary keys of the participating entity sets forms a super key of a relationship set.

- (s_id, i_id) is the super key of *advisor*

- *NOTE: this means **a pair of entities can have at most one relationship in a particular relationship set.***

  ▸ Example: if we wish to track multiple meeting dates between a student and her advisor, we cannot assume a relationship for each meeting. We can use a multivalued attribute though or model meeting as a separate entity

■ Must consider the mapping cardinality of the relationship set when deciding what are the candidate keys

■ Need to consider semantics of relationship set in selecting the *primary key* in case of more than one candidate key

Keys for Relationship Sets Cont.

- Must consider the mapping cardinality of the relationship set when deciding what are the candidate keys

  - 1-1: both primary keys are candidate keys

    - Example: **hasBc**: (Person-Birthcertificate)

  - N-1: the N side is the candidate key

    - Example: **worksFor**: (Instructor-Department)

  - N-M: the combination of both primary keys

    - Example: **takes**: (Student-Course)

Redundant Attributes

- Suppose we have entity sets
  - *instructor*, with attributes including *dept_name*
  - *department*

  and a relationship
  - *inst_dept* relating *instructor* and *department*
- Attribute *dept_name* in entity *instructor* is redundant since there is an explicit relationship *inst_dept* which relates instructors to departments
  - The attribute replicates information present in the relationship, and should be removed from *instructor*
  - BUT: when converting back to tables, in some cases the attribute gets reintroduced, as we will see.

E-R Diagrams

- Rectangles represent entity sets.

- Diamonds represent relationship sets.

- Attributes listed inside entity rectangle

- Underline indicates primary key attributes

Entity With Composite, Multivalued, and Derived Attributes

instructor

<u>ID</u>
name
   first_name
   middle_initial
   last_name
address
   street
      street_number
      street_name
      apt_number
   city
   state
   zip
{ phone_number }
date_of_birth
age ( )
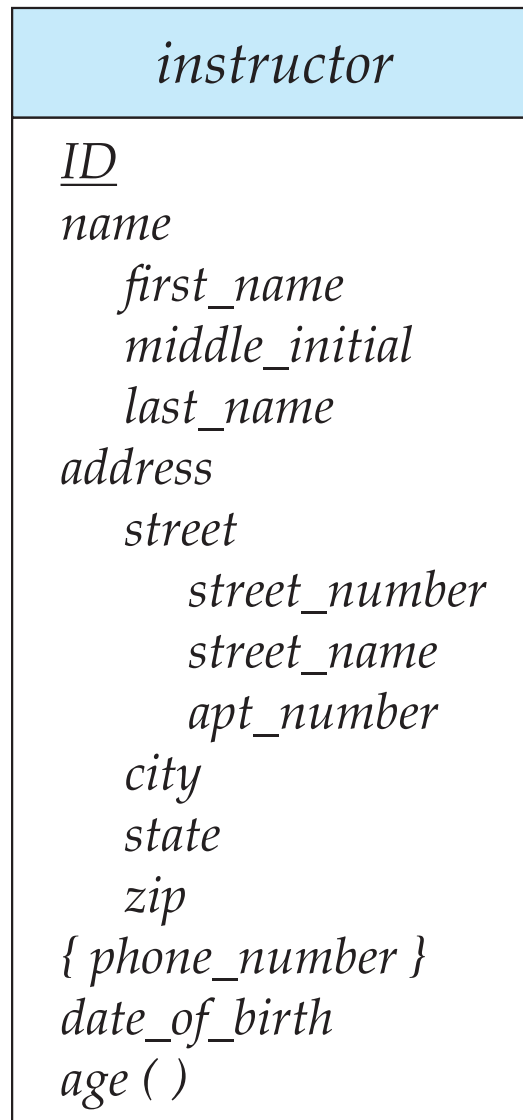
Entity With Composite, Multivalued, and Derived Attributes

Relationship Sets with Attributes

Roles

- Entity sets of a relationship need not be distinct
  - Each occurrence of an entity set plays a "role" in the relationship
- The labels "*course_id*" and "*prereq_id*" are called **roles**.
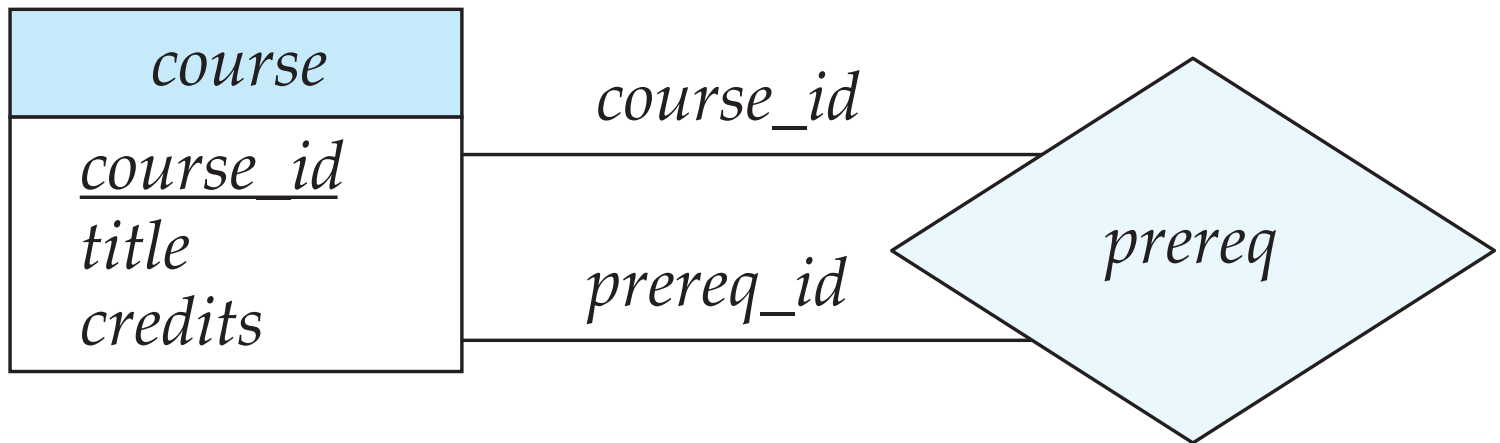
Cardinality Constraints

■ We express cardinality constraints by drawing either a directed line (→), signifying "one," or an undirected line (—), signifying "many," between the relationship set and the entity set.

■ One-to-one relationship:

● A student is associated with at most one *instructor* via the relationship *advisor*

● A *student* is associated with at most one *department* via *stud_dept*

One-to-One Relationship

- one-to-one relationship between an *instructor* and a *student*
  - an instructor is associated with at most one student via *advisor*
  - and a student is associated with at most one instructor via *advisor*

One-to-Many Relationship

- one-to-many relationship between an *instructor* and a *student*
  - an instructor is associated with several (including 0) students via *advisor*
  - a student is associated with at most one instructor via advisor,

Many-to-One Relationships

- In a many-to-one relationship between an *instructor* and a *student,*
  - an instructor is associated with at most one student via *advisor*,
  - and a student is associated with several (including 0) instructors via *advisor*

Many-to-Many Relationship

- An instructor is associated with several (possibly 0) students via *advisor*

- A student is associated with several (possibly 0) instructors via *advisor*

Participation of an Entity Set in a Relationship Set

- **Total participation** (indicated by double line): every entity in the entity set participates in at least one relationship in the relationship set
    - E.g., participation of *section* in *sec_course* is total
        - every *section* must have an associated course
- **Partial participation**: some entities may not participate in any relationship in the relationship set
    - Example: participation of *instructor* in *advisor* is partial

Alternative Notation for Cardinality Limits

■ Cardinality limits can also express participation constraints

Alternative Notation for Cardinality Limits

■ Alternative Notation

E-R Diagram with a Ternary Relationship

Cardinality Constraints on Ternary Relationship

■ We allow at most one arrow out of a ternary (or greater degree) relationship to indicate a cardinality constraint

■ E.g., an arrow from *proj_guide* to *instructor* indicates each student has at most one guide for a project

■ If there is more than one arrow, there are two ways of defining the meaning.

   ● E.g., a ternary relationship *R* between *A*, *B* and *C* with arrows to *B* and *C* could mean

      1. each *A* entity is associated with a unique entity from *B* and *C* or

      2. each pair of entities from (*A, B*) is associated with a unique *C* entity, and each pair (*A, C*) is associated with a unique *B*

   ● Each alternative has been used in different formalisms

   ● To avoid confusion we outlaw more than one arrow

■ Better to use cardinality constraints such as (0,n)

Let's design an ER-model
for
parts of the university database

Partially taken from

Klaus R. Dittrich

modified from:

Database System Concepts, 6th Ed.

Lets design an ER-model for parts of the university database

1) Identify Entities
2) Identify Relationship
3) Determine Attributes
4) Determine Cardinality Constraints

Partially taken from

Klaus R. Dittrich

Weak Entity Sets

- An entity set that does not have a primary key is referred to as a **weak entity set**.

- The existence of a weak entity set depends on the existence of a **identifying entity set**
  - It must relate to the identifying entity set via a total, one-to-many relationship set from the identifying to the weak entity set
  - **Identifying relationship** depicted using a double diamond

- The **discriminator** (*or partial key)* of a weak entity set is the set of attributes that distinguishes among all the entities of a weak entity set that are associated with the same entity of the identifying entity set

- The primary key of a weak entity set is formed by the primary key of the strong entity set on which the weak entity set is existence dependent, plus the weak entity set's discriminator.

Weak Entity Sets (Cont.)

- We underline the discriminator of a weak entity set with a dashed line.

- We put the identifying relationship of a weak entity in a double diamond.

- Primary key for *section* – (*course_id, sec_id, semester, year*)

Weak Entity Sets (Cont.)

- Note: the primary key of the strong entity set is not explicitly stored with the weak entity set, since it is implicit in the identifying relationship.

- If *course_id* were explicitly stored, *section* could be made a strong entity, but then the relationship between *section* and *course* would be duplicated by an implicit relationship defined by the attribute *course_id* common to *course* and *section*

E-R Diagram for a University Enterprise

Reduction to Relational Schemas

Reduction to Relation Schemas

- Entity sets and relationship sets can be expressed uniformly as *relation schemas* that represent the contents of the database.

- A database which conforms to an E-R diagram can be represented by a collection of relation schemas.

- For each entity set and relationship set there is a unique relation schema that is assigned the name of the corresponding entity set or relationship set.

Representing Entity Sets With Simple Attributes

- A strong entity set reduces to a schema with the same attributes
  student(ID, name, tot_cred)

- A weak entity set becomes a table that includes a column for the primary key of the identifying strong entity set
  section (course_id, sec_id, sem, year)

Representing Relationship Sets

- A many-to-many relationship set is represented as a schema with attributes for the primary keys of the two participating entity sets, and any descriptive attributes of the relationship set.

- Example: schema for relationship set *advisor*

advisor = (*s_id, i_id*)

Redundancy of Schemas

- Many-to-one and one-to-many relationship sets that are total on the many-side can be represented by adding an extra attribute to the "many" side, containing the primary key of the "one" side

- Example: Instead of creating a schema for relationship set *inst_dept*, add an attribute *dept_name* to the schema arising from entity set *instructor*
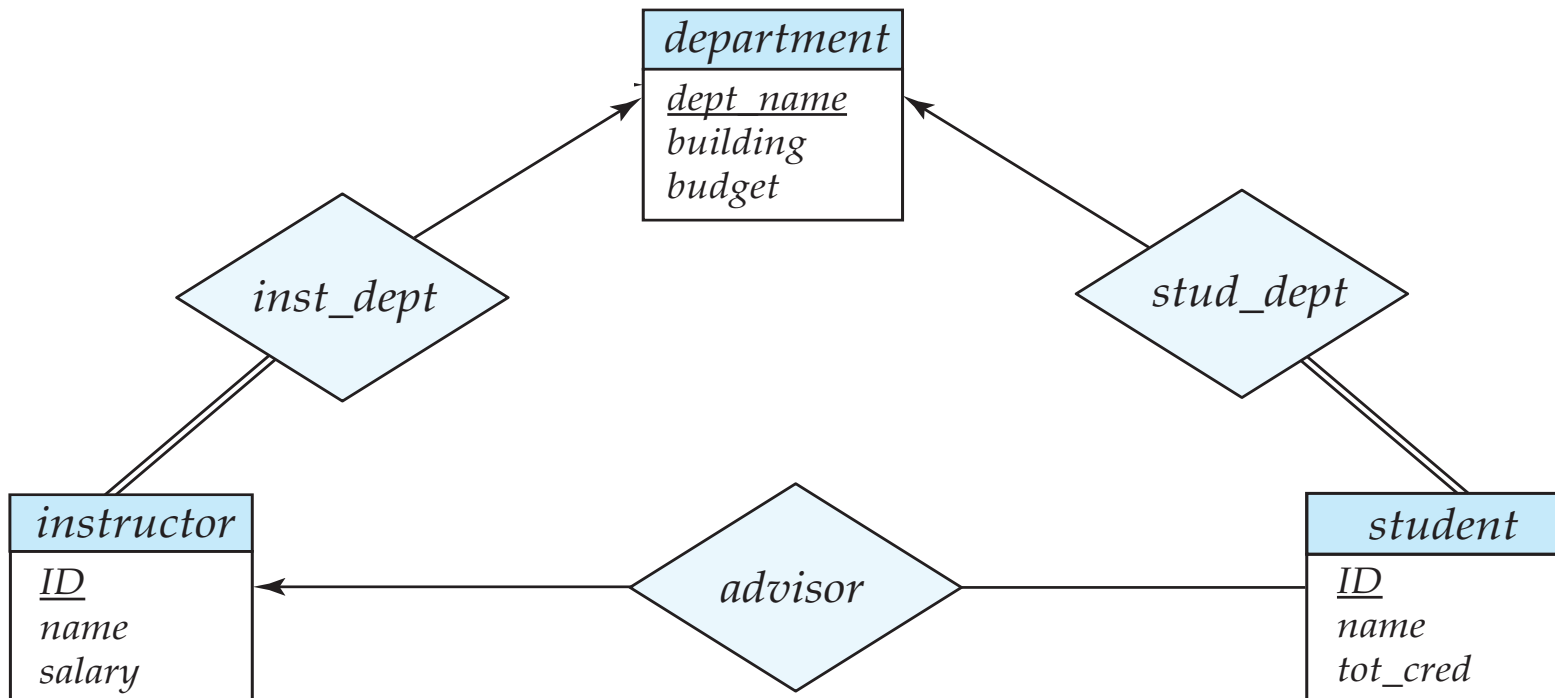
Redundancy of Schemas (Cont.)

- For one-to-one relationship sets, either side can be chosen to act as the "many" side
  - That is, extra attribute can be added to either of the tables corresponding to the two entity sets
  - If the relationship is total in both sides, the relation schemas from the two sides can be merged into one schema
- If participation is *partial* on the "many" side, replacing a schema by an extra attribute in the schema corresponding to the "many" side could result in null values
- The schema corresponding to a relationship set linking a weak entity set to its identifying strong entity set is redundant.
  - Example: The *section* schema already contains the attributes that would appear in the *sec_course* schema

Composite and Multivalued Attributes

instructor

<u>ID</u>
name
   first_name
   middle_initial
   last_name
address
   street
      street_number
      street_name
      apt_number
   city
   state
   zip
{ phone_number }
date_of_birth
age ( )

- Composite attributes are flattened out by creating a separate attribute for each component attribute

  - Example: given entity set *instructor* with composite attribute *name* with component attributes *first_name* and *last_name* the schema corresponding to the entity set has two attributes *name_first_name* and *name_last_name*

    - *Prefix omitted if there is no ambiguity*

- Ignoring multivalued attributes, extended instructor schema is

  - *instructor(ID,*
    first_name, middle_initial, last_name,
    street_number, street_name,
    apt_number, city, state, zip_code,
    date_of_birth)

Composite and Multivalued Attributes

- A multivalued attribute *M* of an entity *E* is represented by a separate schema *EM*

  - Schema *EM* has attributes corresponding to the primary key of *E* and an attribute corresponding to multivalued attribute *M*

  - Example: Multivalued attribute *phone_number* of *instructor* is represented by a schema:
    inst_phone= (*ID, phone_number*)

  - Each value of the multivalued attribute maps to a separate tuple of the relation on schema *EM*

    ▸ For example, an *instructor* entity with primary key 22222 and phone numbers 456-7890 and 123-4567 maps to two tuples:
      (22222, 456-7890) and (22222, 123-4567)

Multivalued Attributes (Cont.)

- Special case:entity *time_slot* has only one attribute other than the primary-key attribute, and that attribute is multivalued

  - Optimization: Don't create the relation corresponding to the entity, just create the one corresponding to the multivalued attribute

  - *time_slot*(*time_slot_id, day, start_time, end_time*)

  - Caveat: *time_slot* attribute of *section (*from *sec_time_slot*) cannot be a foreign key due to this optimization

Design Issues

- **Use of entity sets vs. attributes**
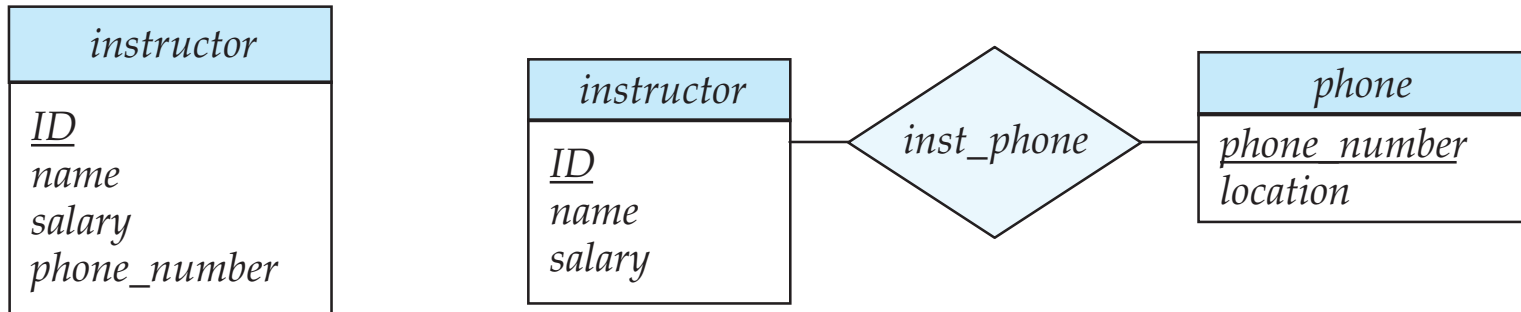
- Designing phone as an entity allow for primary key constraints for phone

- Designing phone as an entity allow phone numbers to be used in relationships with other entities (e.g., student)

- Use of phone as an entity allows extra information about phone numbers

Design Issues

- **Use of entity sets vs. relationship sets**
  - Possible guideline is to designate a relationship set to describe an action that occurs between entities
  - Possible hint: the relationship only relates entities, but does not have an existence by itself. E.g., hasAddress: (department-address)
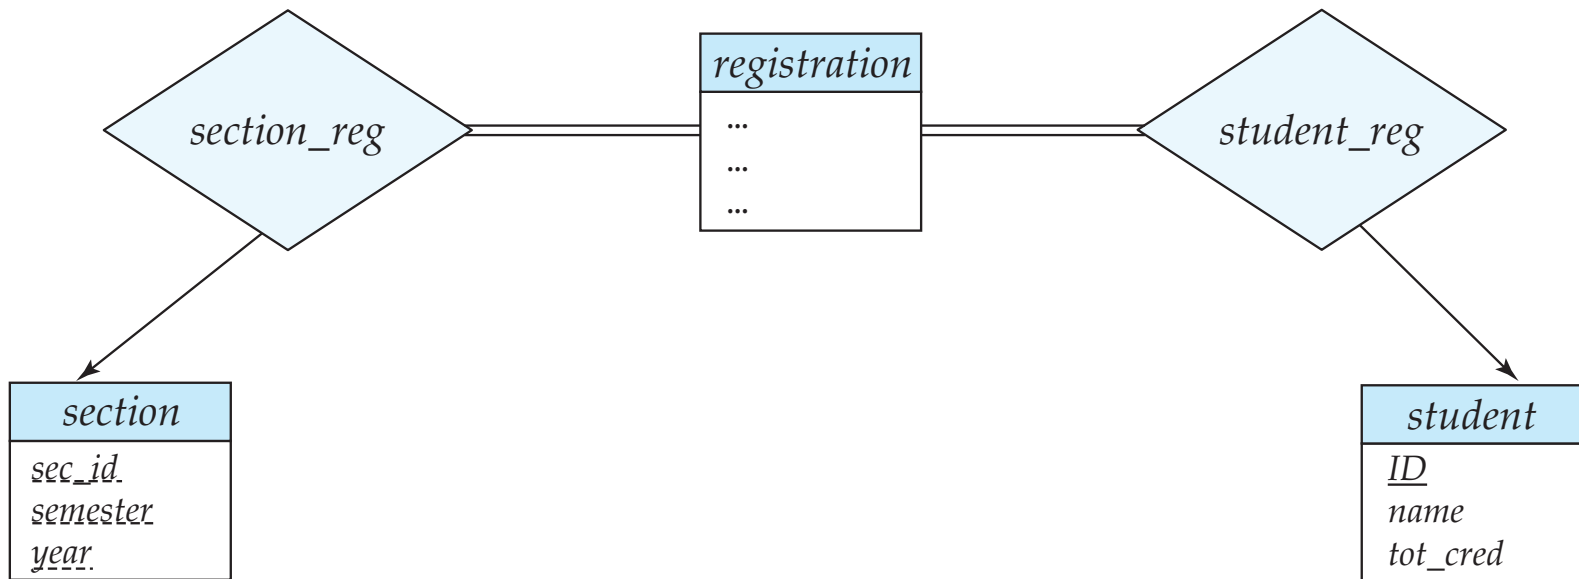
Design Issues

- **Binary versus n-ary relationship sets**

  - Although it is possible to replace any nonbinary (n-ary, for $n > 2$) relationship set by a number of distinct binary relationship sets + an aritifical entity set, a n-ary relationship set shows more clearly that several entities participate in a single relationship.

- **Placement of relationship attributes**

  - e.g., attribute *date* as attribute of *advisor* or as attribute of *student*

  - Does not work for *N-M relationships!*

Binary Vs. Non-Binary Relationships

■ Some relationships that appear to be non-binary may be better represented using binary relationships

- E.g., A ternary relationship *parents*, relating a child to his/her father and mother, is best replaced by two binary relationships, *father* and *mother*

  ▸ Using two binary relationships allows partial information (e.g., only mother being know)

- But there are some relationships that are naturally non-binary

  ▸ Example: *proj_guide*

Converting Non-Binary Relationships to Binary Form

- In general, any non-binary relationship can be represented using binary relationships by creating an artificial entity set.

  - Replace R between entity sets A, B and C by an entity set E, and three relationship sets:

    1. R_A, relating E and A 2. R_B, relating E and B
    3. R_C, relating E and C

  - Create a special identifying attribute for E

  - Add any attributes of R to E

  - For each relationship (a_i, b_i, c_i) in R, create

    1. a new entity e_i in the entity set E 2. add (e_i, a_i) to R_A

    3. add (e_i, b_i) to R_B 4. add (e_i, c_i) to R_C

(a)   (b)

Converting Non-Binary Relationships (Cont.)

- **Also need to translate constraints**
  - Translating all constraints may not be possible
  - There may be instances in the translated schema that cannot correspond to any instance of R
    - Exercise: *add constraints to the relationships R_A, R_B and R_C to ensure that a newly created entity corresponds to exactly one entity in each of entity sets A, B and C*
  - We can avoid creating an identifying attribute by making E a weak entity set (described shortly) identified by the three relationship sets

Converting Non-Binary Relationships: Is the New Entity Set E Necessary?

- Yes, because a non-binary relation ship stores more information that any number of binary relationships

  - Consider again the example (a) below

  - Replace R with three binary relationships:

    1. R_{AB}, relating A and B   2. R_{BC}, relating B and C
    3. R_{AC}, relating A and C

  - For each relationship (a_i , b_i , c_i) in R, create

    ▸ 1. add (a_i , b_i) to R_{AB}

    ▸ 2. add (b_i , c_i) to R_{BC}

    ▸ 3. add (a_i , c_i) to R_{AC}

  - *Consider R = order, A = supplier, B = item, C = customer*

(Gunnar, chainsaw, Bob) – Bob ordered a chainsaw from Gunnar

->

(Gunnar, chainsaw), (chainsaw, Bob), (Gunnar, Bob)

Gunnar supplies chainsaws, Bob ordered a chainsaw, Bob ordered something from Gunnar. E.g., we do not know what Bob ordered from Gunnar.

ER-model to Relational Summary

■ ***Rule 1)*** *Strong entity E*

- *Create relation with attributes of E*
- *Primary key is equal to the PK of E*

■ ***Rule 2)*** *Weak entity W identified by E through relationship R*

- *Create relation with attributes of W and R and PK(E).*
- *Set PK to discriminator attributes combined with PK(E). PK(E) is a foreign key to E.*

■ ***Rule 3)*** *Binary relationship R between A and B: one-to-one*

- *If no side is total add PK of A to as foreign key in B or the other way around. Add any attributes of the relationship R to A respective B.*
- *If one side is total add PK of the other-side as foreign key. Add any attributes of the relationship R to the total side.*
- *If both sides are total merge the two relation into a new relation E and choose either PK(A) as PK(B) as the new PK. Add any attributes of the relationship R to the new relation E.*

ER-model to Relational Summary (Cont.)

- **■** *Rule 4)* *Binary relationship R between A and B: one-to-many/many-to-one*
  - **●** *Add PK of the "one" side as foreign key to the "many" side.*
  - **●** *Add any attributes of the relationship R to the "many" side.*
- **■** *Rule 5)* *Binary relationship R between A and B: many-to-many*
  - **●** *Create a new relation R.*
  - **●** *Add PK's of A and B as attributes + plus all attributes of R.*
  - **●** *The primary key of the relationship is PK(A) + PK(B). The PK attributes of A/B form a foreign key to A/B*
- **■** *Rule 6)* *N-ary relationship R between $E_1 \ldots E_n$*
  - **●** *Create a new relation.*
  - **●** *Add all the PK's of $E_1 \ldots E_n$. Add all attributes of R to the new relation.*
  - **●** *The primary key or R is $PK(E_1) \ldots PK(E_n)$. Each $PK(E_i)$ is a foreign key to the corresponding relation.*

ER-model to Relational Summary (Cont.)

- **Rule 7)** *Entity E with multi-valued attribute A*
    - *Create new relation. Add A and PK(E) as attributes.*
    - *PK is all attributes. PK(E) is a foreign key.*

E-R Diagram for a University Enterprise
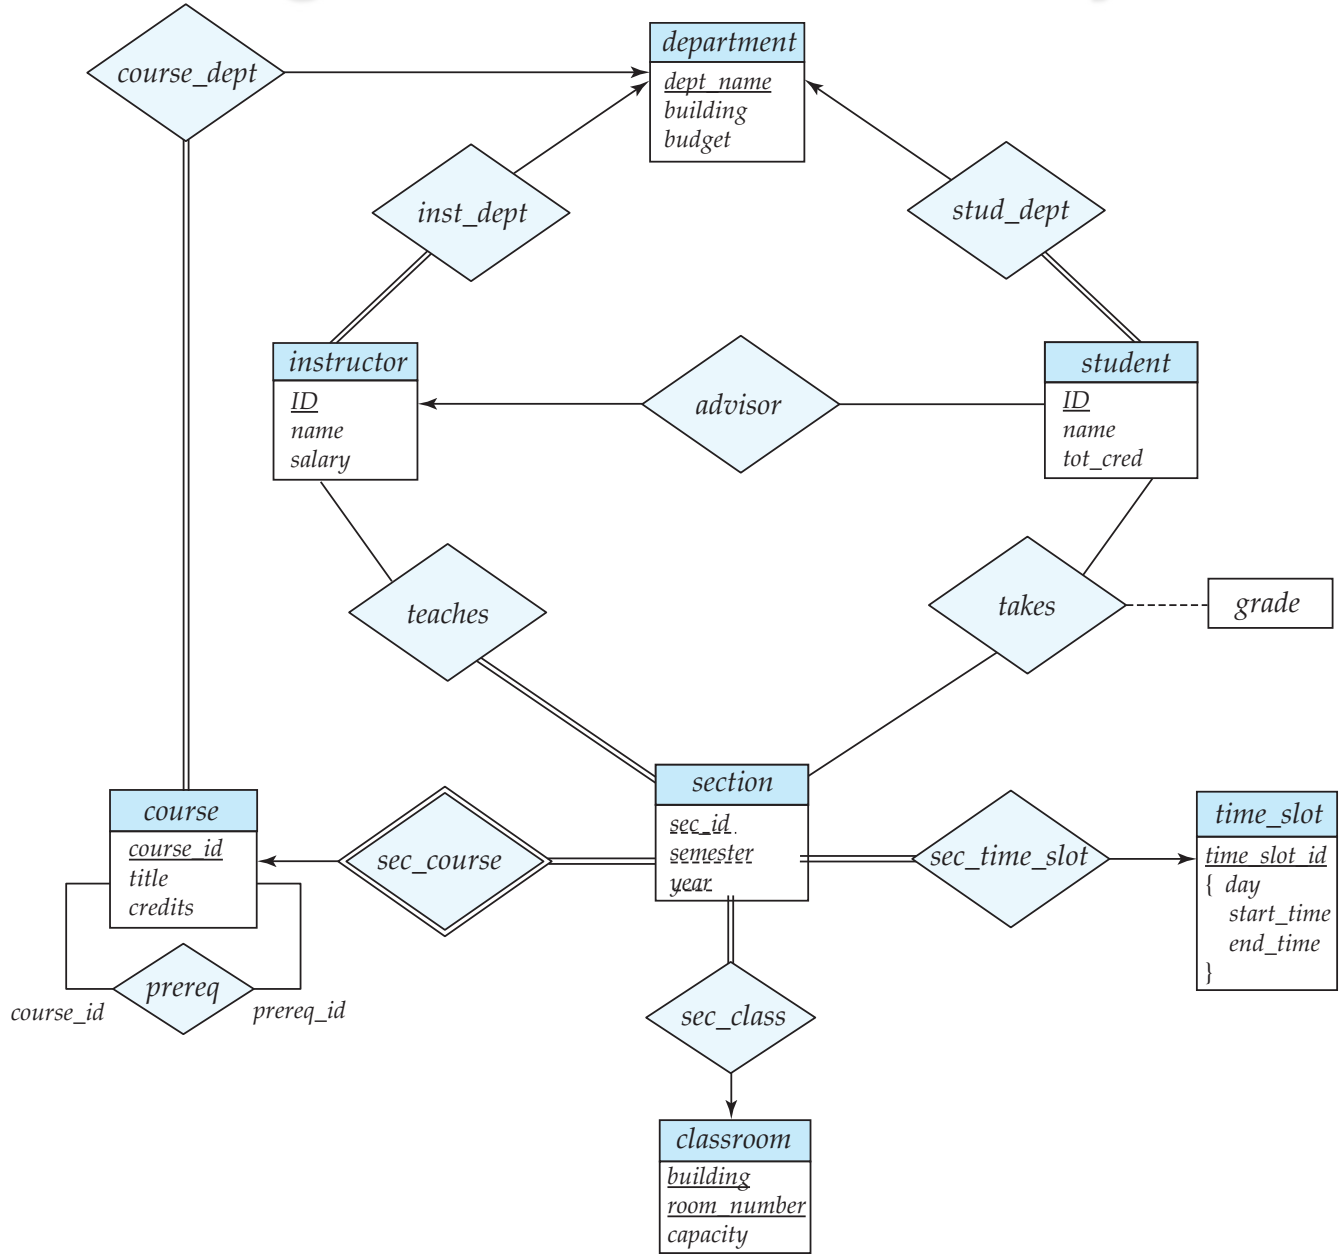
Translate the University ER-Model

- **Rule 1) Strong Entities**

  - **department**(<u>dept_name</u>, building, budget)

  - **instructor**(<u>ID</u>, name, salary)

  - **student**(<u>ID</u>, name, tot_cred)

  - **course**(<u>course_id</u>, title, credits)

  - **time_slot**(<u>time_slot_id</u>)

  - **classroom**(<u>building</u>,<u>room_number</u>, capacity)

- **Rule 2) Weak Entities**

  - **section**(<u>course_id</u>, <u>sec_id</u>, <u>semester</u>, <u>year</u>)
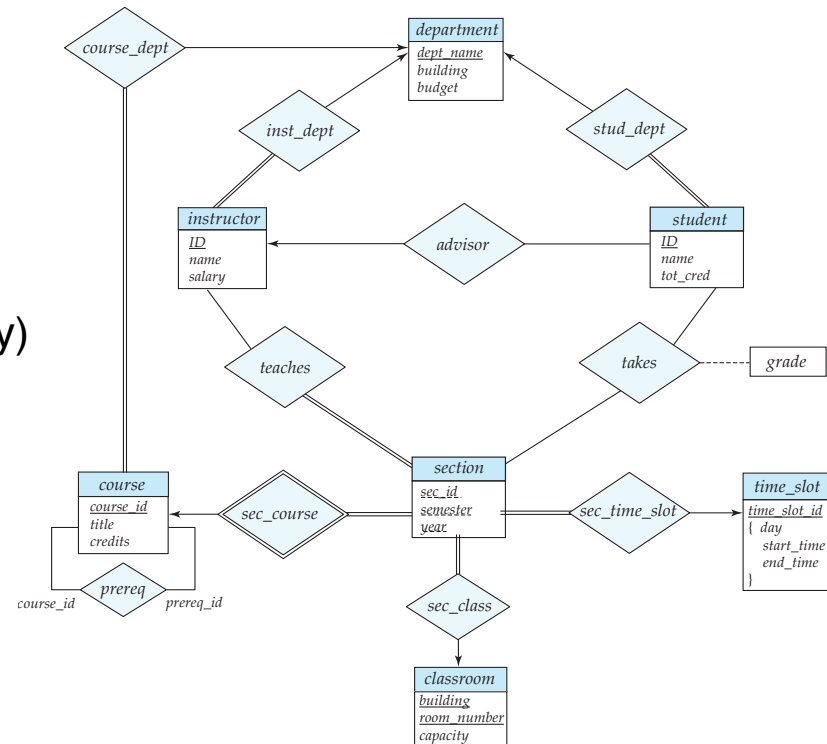
Translate the University ER-Model

- **Rule 3) Relationships one-to-one**
  - None exist

- **Rule 4) Relationships one-to-many**
  - **department**(<u>dept_name</u>, building, budget)
  - **instructor**(<u>ID</u>, name, salary, **dept_name**)
  - **student**(<u>ID</u>, name, tot_cred, **dept_name**, **instr_ID**)
  - **course**(<u>course_id</u>, title, credits, **dept_name**)
  - **time_slot**(<u>time_slot_id</u>)
  - **classroom**(<u>building</u>,<u>room_number</u>, capacity)
  - **section**(<u>course_id</u>, <u>sec_id</u>, <u>semester</u>, <u>year</u>, **room_building**, **room_number**, **time_slot_id**)
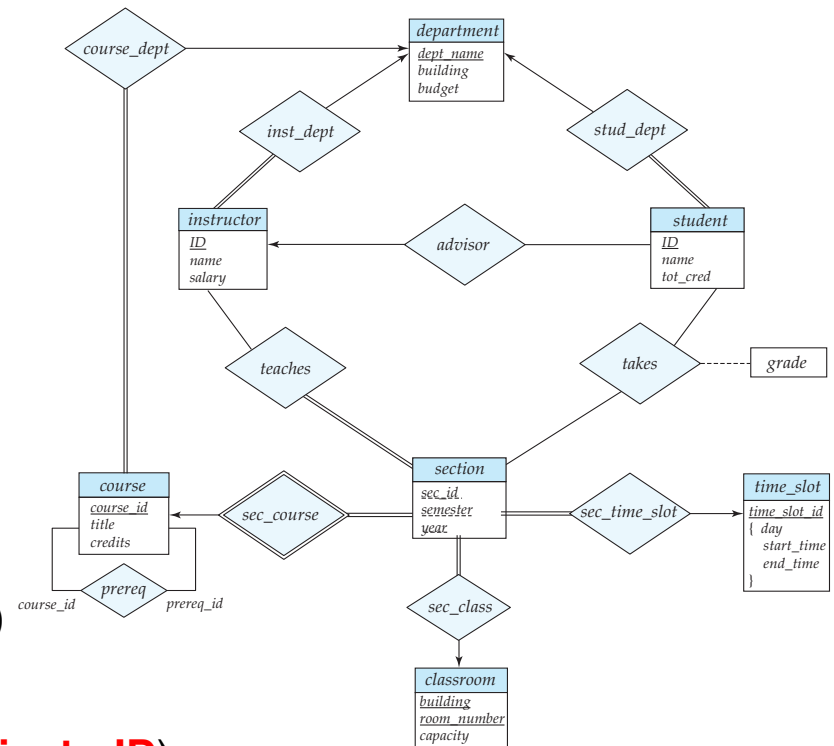
Translate the University ER-Model

- **Rule 5) Relationships many-to-many**
  - **department**(<u>dept_name</u>, building, budget)
  - **instructor**(<u>ID</u>, name, salary, dept_name)
  - **student**(<u>ID</u>, name, tot_cred, dept_name, instr_ID)
  - **course**(<u>course_id</u>, title, credits, dept_name)
  - **time_slot**(<u>time_slot_id</u>)
  - **classroom**(<u>building,room_number</u>, capacity)
  - **section**(<u>course_id</u>, <u>sec_id</u>, <u>semester</u>, <u>year</u>, room_building, room_number, time_slot_id)
  - **prereq**(<u>course_id</u>, <u>prereq_id</u>)
  - **teaches**(<u>ID</u>, <u>course_id</u>, <u>sec_id</u>, <u>semester</u>, <u>year</u>)
  - **takes**(<u>ID</u>, <u>course_id</u>, <u>sec_id</u>, <u>semester</u>, <u>year</u>, grade)
- **Rule 6) N-ary Relationships**
  - none exist
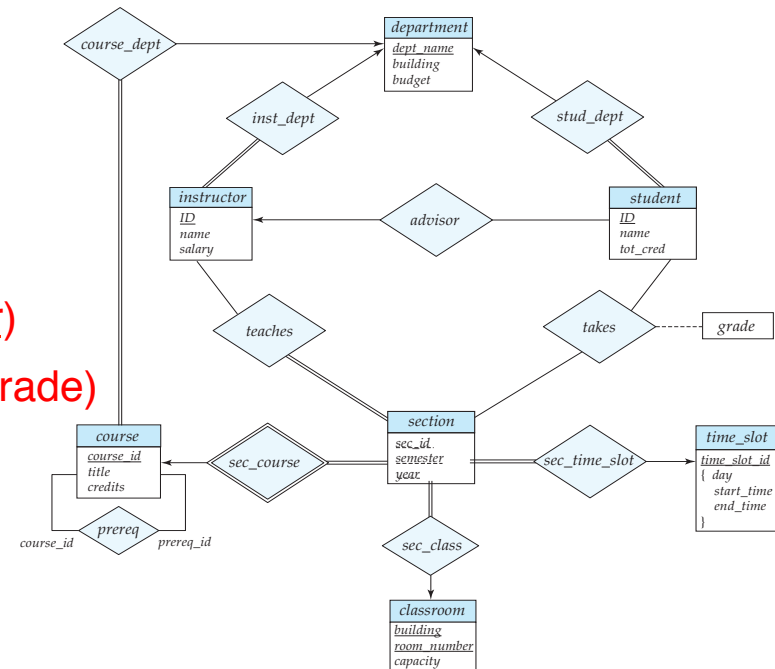
Translate the University ER-Model

- **Rule 7) Multivalued attributes**

  - **department**(<u>dept_name</u>, building, budget)

  - **instructor**(<u>ID</u>, name, salary, dept_name)

  - **student**(<u>ID</u>, name, tot_cred, dept_name, instr_ID)

  - **course**(<u>course_id</u>, title, credits, dept_name)

  - **time_slot**(time_slot_id)

  - **time_slot_day**(<u>time_slot_id</u>, <u>start_time</u>, <u>end_time</u>)

  - **classroom**(<u>building</u>,<u>room_number</u>, capacity)

  - **section**(<u>course_id</u>, <u>sec_id</u>, <u>semester</u>, <u>year</u>, room_building, room_number, time_slot_id)

  - **prereq**(<u>course_id</u>, <u>prereq_id</u>)

  - **teaches**(<u>ID</u>, <u>course_id</u>, <u>sec_id</u>, <u>semester</u>, <u>year</u>)

  - **takes**(<u>ID</u>, <u>course_id</u>, <u>sec_id</u>, <u>semester</u>, <u>year</u>, grade)
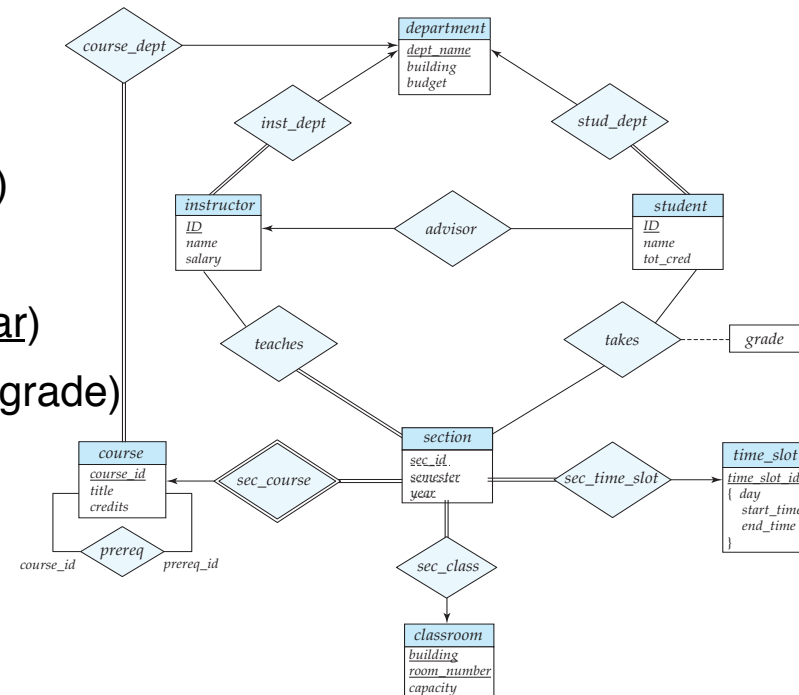
Extended ER Features

Extended E-R Features: Specialization

- Top-down design process; we designate subgroupings within an entity set that are distinctive from other entities in the set.

- These subgroupings become lower-level entity sets that have attributes or participate in relationships that do not apply to the higher-level entity set.

- Depicted by a *triangle* component labeled ISA (E.g., *instructor* "is a" *person*).

- **Attribute inheritance** – a lower-level entity set inherits all the attributes and relationship participation of the higher-level entity set to which it is linked.

Specialization Example
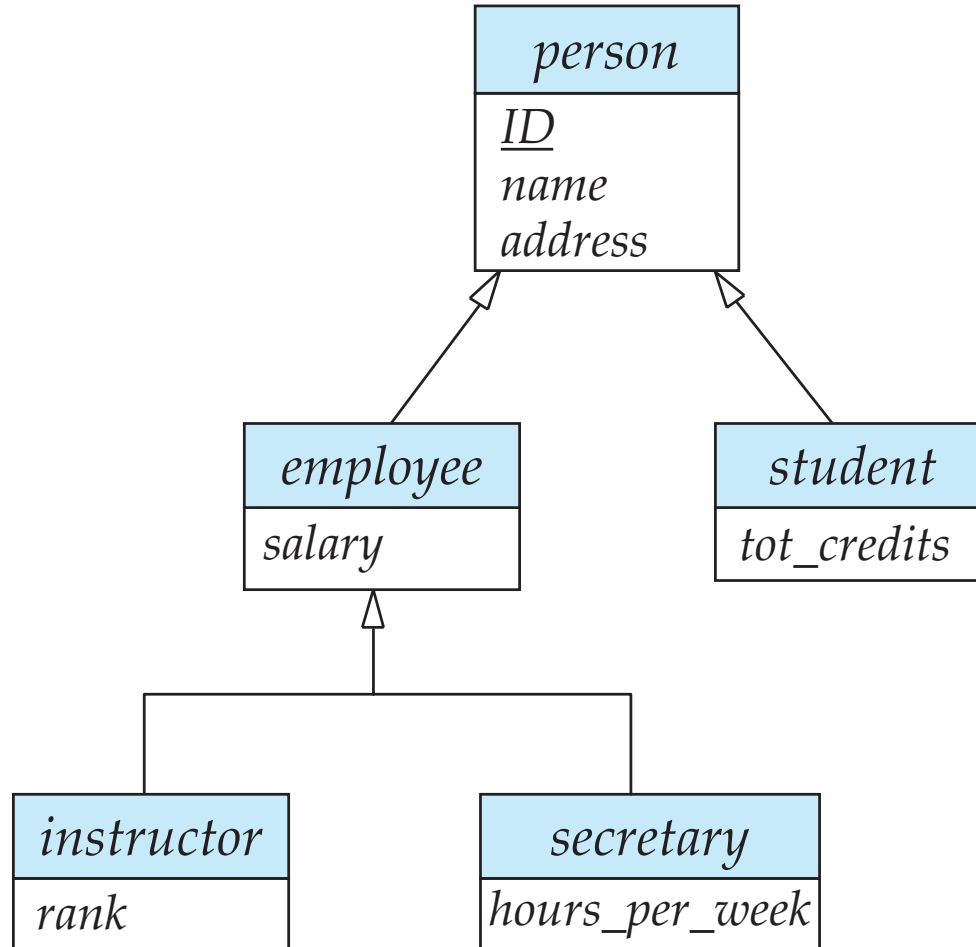
Extended ER Features: Generalization

- **A bottom-up design process** – combine a number of entity sets that share the same features into a higher-level entity set.

- Specialization and generalization are simple inversions of each other; they are represented in an E-R diagram in the same way.

- The terms specialization and generalization are used interchangeably.

Specialization and Generalization (Cont.)

- Can have multiple specializations of an entity set based on different features.

- E.g., *permanent_employee* vs. *temporary_employee*, in addition to *instructor* vs. *secretary*

- Each particular employee would be

  - a member of one of *permanent_employee* or *temporary_employee*,

  - and also a member of one of *instructor*, *secretary*

- The ISA relationship also referred to as **superclass - subclass** relationship

Design Constraints on a Specialization/ Generalization

- Constraint on which entities can be members of a given lower-level entity set.

  - condition-defined

    - Example: all customers over 65 years are members of *senior-citizen* entity set; *senior-citizen* ISA *person*.

  - user-defined

- Constraint on whether or not entities may belong to more than one lower-level entity set within a single generalization.

  - **Disjoint**

    - an entity can belong to only one lower-level entity set

    - Noted in E-R diagram by having multiple lower-level entity sets link to the same triangle

  - **Overlapping**

    - an entity can belong to more than one lower-level entity set

Specialization Example

Disjoint, employees are either instructors or secretaries

person
ID
name
address
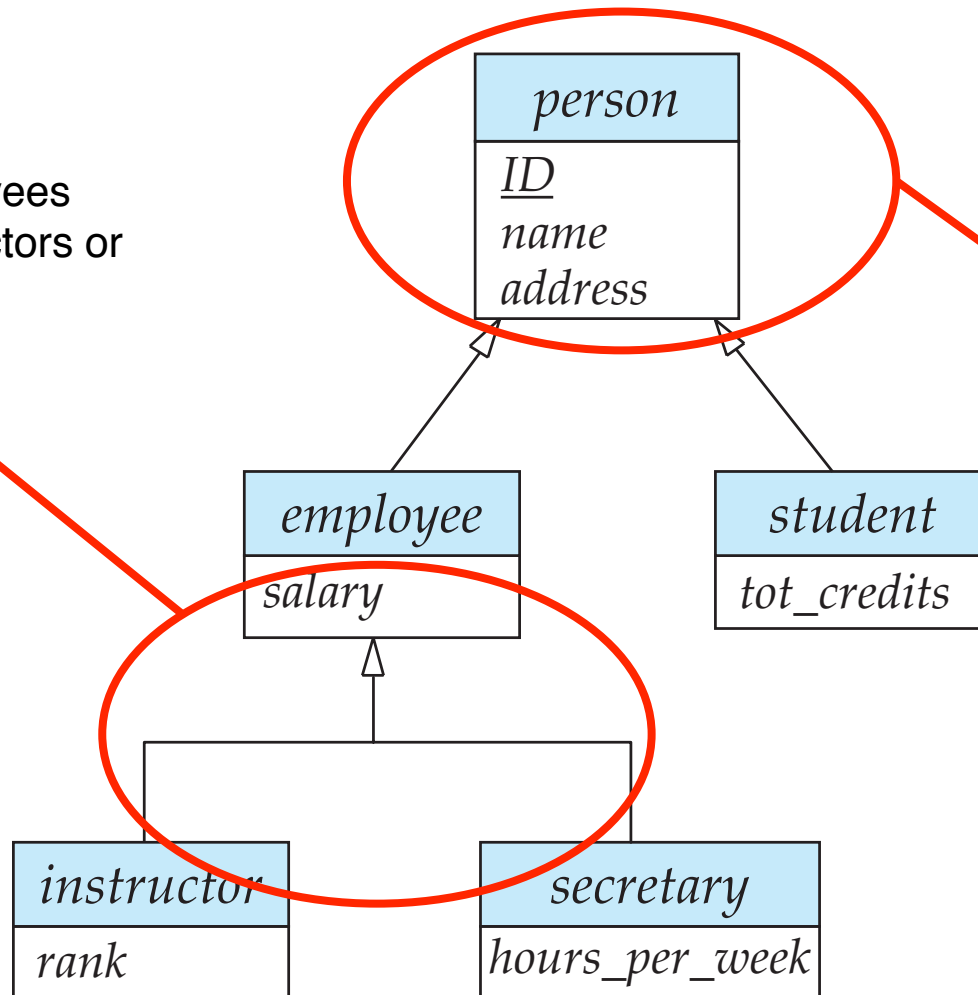
Overlapping, a person can be both an employee and a student

employee
salary

student
tot_credits

instructor
rank

secretary
hours_per_week

Design Constraints on a Specialization/ Generalization (Cont.)

■ **Completeness constraint** -- specifies whether or not an entity in the higher-level entity set must belong to at least one of the lower-level entity sets within a generalization.

- ● **total**: an entity must belong to one of the lower-level entity sets

- ● **partial**: an entity need not belong to one of the lower-level entity sets

Aggregation

■ Consider the ternary relationship *proj_guide*, which we saw earlier

■ Suppose we want to record evaluations of a student by a guide on a project
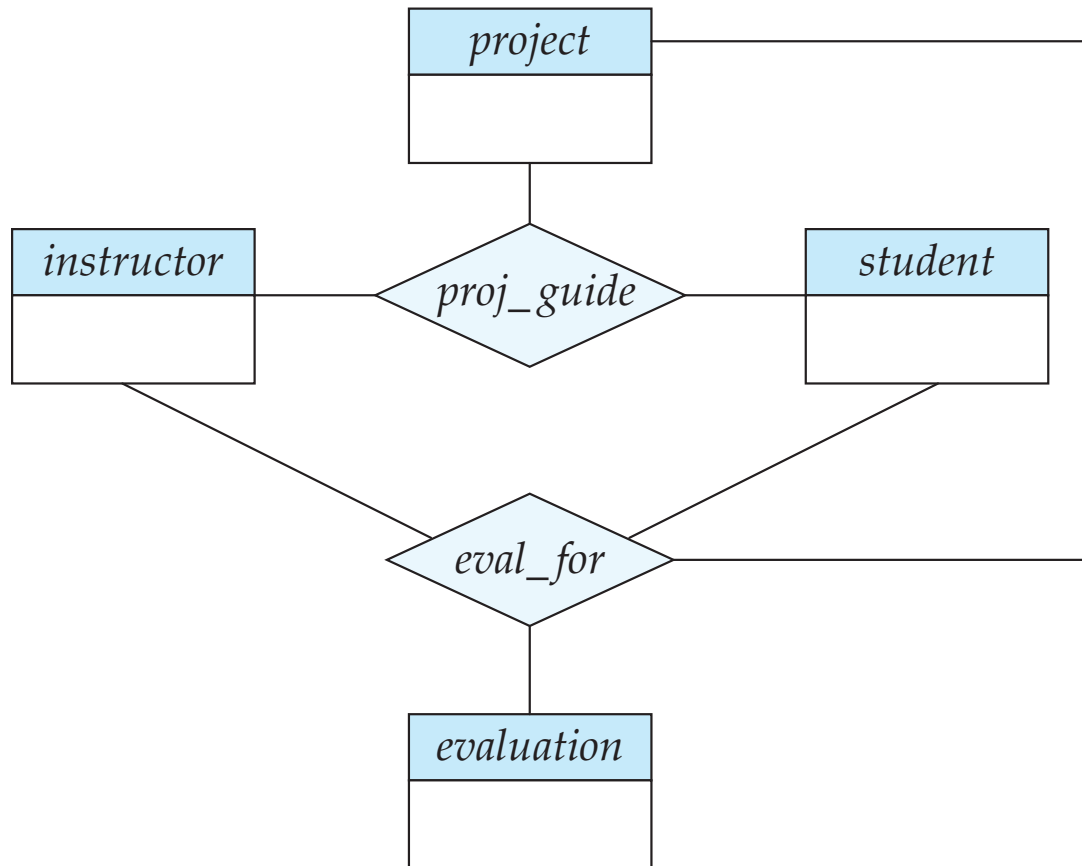
Aggregation (Cont.)

- Relationship sets *eval_for* and *proj_guide* represent overlapping information

    - Every *eval_for* relationship corresponds to a *proj_guide* relationship

    - However, some *proj_guide* relationships may not correspond to any *eval_for* relationships

        ▸ So we can't discard the *proj_guide* relationship

- Eliminate this redundancy via *aggregation*

    - Treat relationship as an abstract entity

    - Allows relationships between relationships

    - Abstraction of relationship into new entity

Aggregation (Cont.)

- Without introducing redundancy, the following diagram represents:
  - A student is guided by a particular instructor on a particular project
  - A student, instructor, project combination may have an associated evaluation

Representing Specialization via Schemas

- Method 1:
  - Form a relation schema for the higher-level entity
  - Form a relation schema for each lower-level entity set, include primary key of higher-level entity set and local attributes

schema	attributes
person	*ID, name, street, city*
student	*ID, tot_cred*
employee	*ID, salary*

  - Drawback: getting information about, an *employee* requires accessing two relations, the one corresponding to the low-level schema and the one corresponding to the high-level schema

Representing Specialization as Schemas (Cont.)

- Method 2:
  - Form a single relation schema for each entity set with all local and inherited attributes

schema	attributes
person	*ID, name, street, city*
student	*ID, name, street, city, tot_cred*
employee	*ID, name, street, city, salary*

  - If specialization is total, the schema for the generalized entity set (*person*) not required to store information
    - Can be defined as a "view" relation containing union of specialization relations
    - But explicit schema may still be needed for foreign key constraints
  - Drawback: *name, street* and *city* may be stored redundantly for people who are both students and employees

Representing Specialization as Schemas (Cont.)

- Method 3:
  - Form a single relation schema for each entity set with all local and inherited attributes
    - ▸ For total and disjoint specialization add a single "type" attribute that stores the type of an entity

schema	attributes
person	*ID, **type**, name, street, city, tot_cred, salary*

  - ▸ For partial and/or overlapping specialization add multiple boolean "type" attributes

schema	attributes
person	*ID, **isEmployee, isStudent**, name, street, city, tot_cred, salary*

  - Drawback: large number of NULL values, potentially large relation

Schemas Corresponding to Aggregation

■ To represent aggregation, create a schema containing

- primary key of the aggregated relationship,
- the primary key of the associated entity set
- any descriptive attributes

- For example, to represent aggregation manages between relationship works_on and entity set manager, create a schema

  eval_for (*s_ID, project_id, i_ID, evaluation_id*)

ER-model to Relational Summary (Cont.)

- **Rule 8)** *Specialization of E into $S_1, \ldots, S_n$ (method 1)*
  - *Create a relation for E with all attributes of E. The PK of E is the PK.*
  - *For each S_i create a relation with PK(E) as PK and foreign key to relation for E. Add all attributes of S_i that do not exist in E.*
- **Rule 9)** *Specialization of E into $S_1, \ldots, S_n$ (method 2)*
  - *Create a relation for E with all attributes of E. The PK of E is the PK.*
  - *For each S_i create a relation with PK(E) as PK and foreign key to relation for E. Add all attributes of S_i.*
- **Rule 10)** *Specialization of E into $S_1, \ldots, S_n$ (method 3)*
  - *Create a new relation with all attributes from E and $S_1, \ldots, S_n$.*
  - *Add single attribute type or a boolean type attribute for each S_i*
  - *The primary key is PK(E)*

■ **Rule 11)** *Aggregation: Relationship R_1 relates entity sets E_1, ..., E_n. This is related by relationship A to an entity set B*

- *Create a relation for A with attributes $PK(E_1)$ ... $PK(E_n)$ + all attributes from A + PK(B). PK are all attributes except the ones from A*

ER Design Decisions

- The use of an attribute or entity set to represent an object.

- Whether a real-world concept is best expressed by an entity set or a relationship set.

- The use of a ternary relationship versus a pair of binary relationships.

- The use of a strong or weak entity set.

- The use of specialization/generalization – contributes to modularity in the design.

- The use of aggregation – can treat the aggregate entity set as a single unit without concern for the details of its internal structure.

How about doing another ER design interactively on the board?

Partially taken from

Klaus R. Dittrich

modified from:

Database System Concepts, 6th Ed.

Summary of Symbols Used in E-R Notation

E   entity set

R   relationship set

R   identifying relationship set for weak entity set

R ─ E   total participation of entity set in relationship

E
A1
A2
A2.1
A2.2
{A3}
A4()

attributes:
simple (A1),
composite (A2) and
multivalued (A3)
derived (A4)

E
A1

primary key

E
A1

discriminating attribute of weak entity set

Symbols Used in ER Notation (Cont.)

Alternative ER Notations

- Chen, IDE1FX, …

entity set E with
simple attribute A1,
composite attribute A2,
multivalued attribute A3,
derived attribute A4,
and primary key A1

weak entity set   generalization ISA   total generalization ISA

Alternative ER Notations

UML

- **UML**: Unified Modeling Language

- UML has many components to graphically model different aspects of an entire software system

- UML Class Diagrams correspond to E-R Diagram, but several differences.

ER vs. UML Class Diagrams

ER Diagram Notation

Equivalent in UML

E
A1
M1()

entity with attributes (simple, composite, multivalued, derived)

E
−A1
+M1()

class with simple attributes and methods (attribute prefixes: + = public, − = private, # = protected)

E1 —role1— ⟨R⟩ —role2— E2   binary relationship

E1 —role1— R —role2— E2

A1 ⋯ ⟨R⟩   E1 —role1— ⟨R⟩ —role2— E2   relationship attributes

R
A1

E1 —role1— ⋮ —role2— E2

E1 —0..*— ⟨R⟩ —0..1— E2   cardinality constraints

E1 —0..1— R —0..*— E2

*Note reversal of position in cardinality constraint depiction

ER vs. UML Class Diagrams

ER Diagram Notation

Equivalent in UML

n-ary relationships

overlapping generalization

overlapping

disjoint generalization

disjoint

*Generalization can use merged or separate arrows independent of disjoint/overlapping

UML Class Diagrams (Cont.)

- Binary relationship sets are represented in UML by just drawing a line connecting the entity sets. The relationship set name is written adjacent to the line.

- The role played by an entity set in a relationship set may also be specified by writing the role name on the line, adjacent to the entity set.

- The relationship set name may alternatively be written in a box, along with attributes of the relationship set, and the box is connected, using a dotted line, to the line depicting the relationship set.

Recap

- ER-model
  - Entities
    - Strong
    - Weak
  - Attributes
    - Simple vs. Composite
    - Single-valued vs. Multi-valued
  - Relationships
    - Degree (binary vs. N-ary)
  - Cardinality constraints
  - Specialization/Generalization
    - Total vs. partial
    - Disjoint vs. overlapping
  - Aggregation

Recap Cont.

- ER-Diagrams
  - Alternative notations
- UML-Diagrams
- Design decisions
  - Multi-valued attribute vs. entity
  - Entity vs. relationship
  - Binary vs. N-ary relationships
  - Placement of relationship attributes
  - Total 1-1 vs. single entity
- ER to relational model
  - Translation rules

End of Chapter 7

Partially taken from

Klaus R. Dittrich

modified from:

Database System Concepts, 6th Ed.

Outline

- Introduction
- Relational Data Model
- Formal Relational Languages (relational algebra)
- SQL - Advanced
- **Database Design – Database modelling**
- Transaction Processing, Recovery, and Concurrency Control
- Storage and File Structures
- Indexing and Hashing
- Query Processing and Optimization

Figure 7.01

76766	Crick
45565	Katz
10101	Srinivasan
98345	Kim
76543	Singh
22222	Einstein

instructor

98988	Tanaka
12345	Shankar
00128	Zhang
76543	Brown
76653	Aoi
23121	Chavez
44553	Peltier

student

Figure 7.02

instructor

student
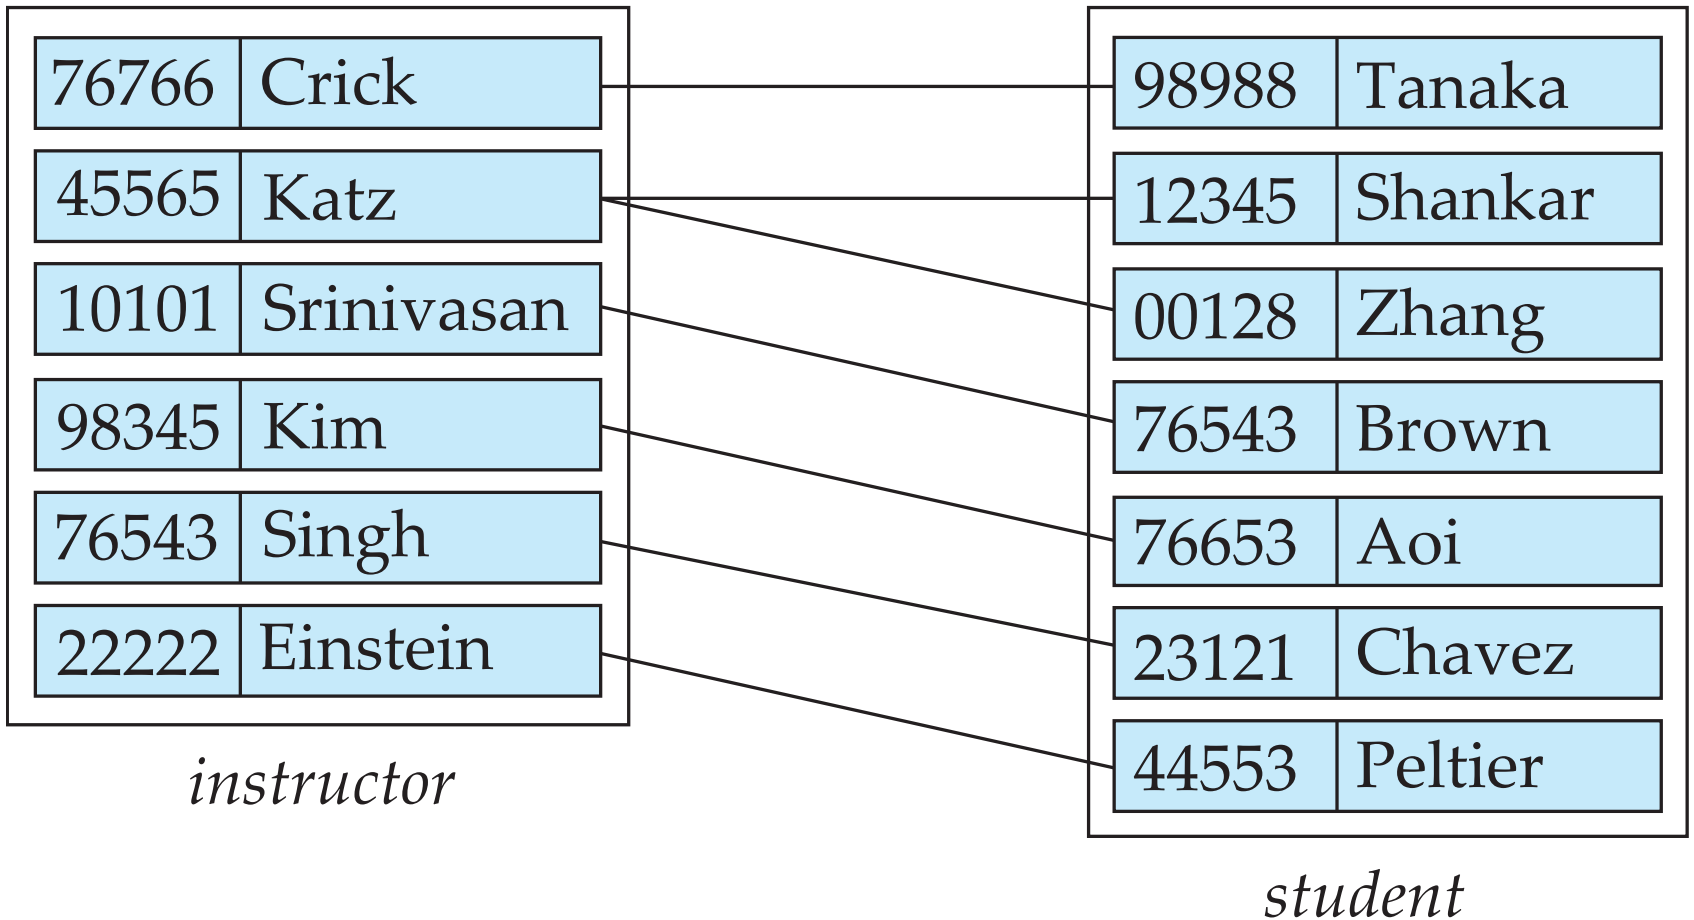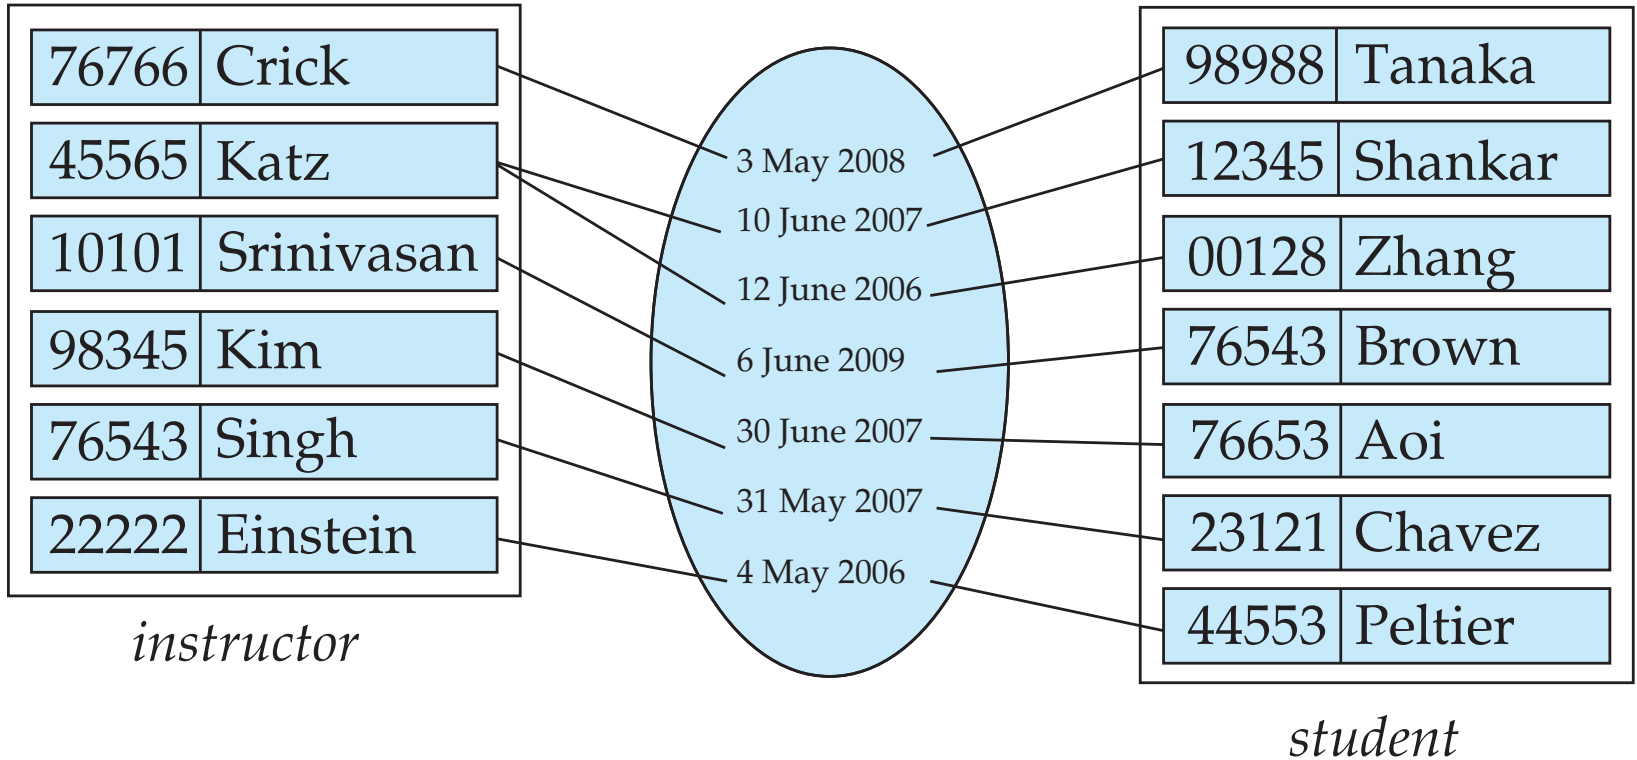
Figure 7.03

Figure 7.04

composite
attributes

name

first_name middle_initial last_name

address

street city state postal_code

component
attributes

street_number street_name apartment_number

Figure 7.05

(a)   (b)

Figure 7.06

(a)   (b)

Figure 7.07

Figure 7.08

Figure 7.09

(a)

(b)

(c)

Figure 7.10

Figure 7.11

instructor
ID
name
first_name
middle_initial
last_name
address
street
street_number
street_name
apt_number
city
state
zip
{ phone_number }
date_of_birth
age ( )

Figure 7.12
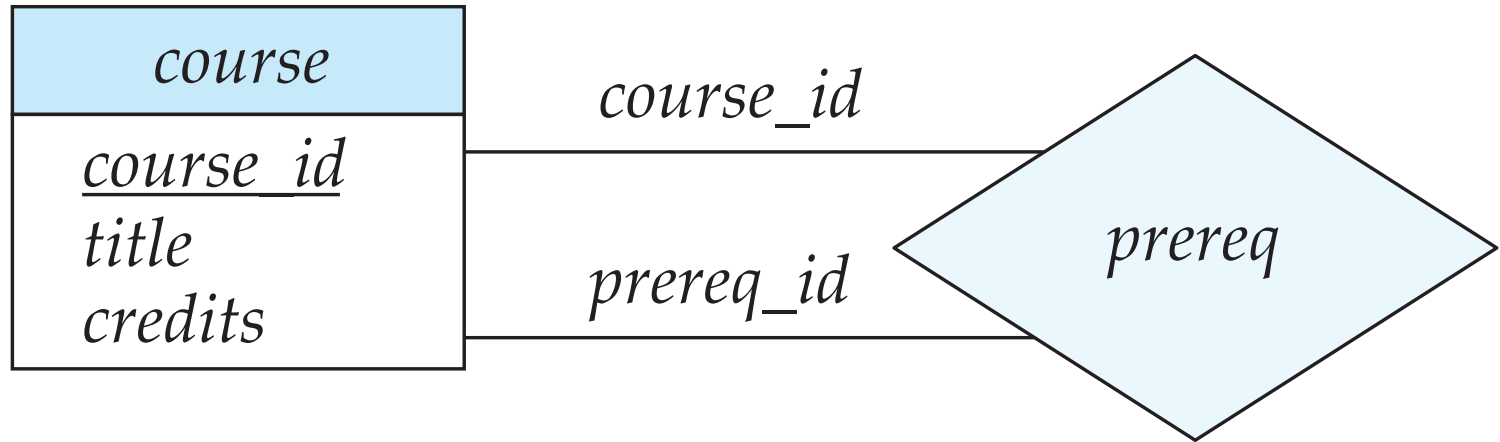
Figure 7.13

Figure 7.14

Figure 7.15
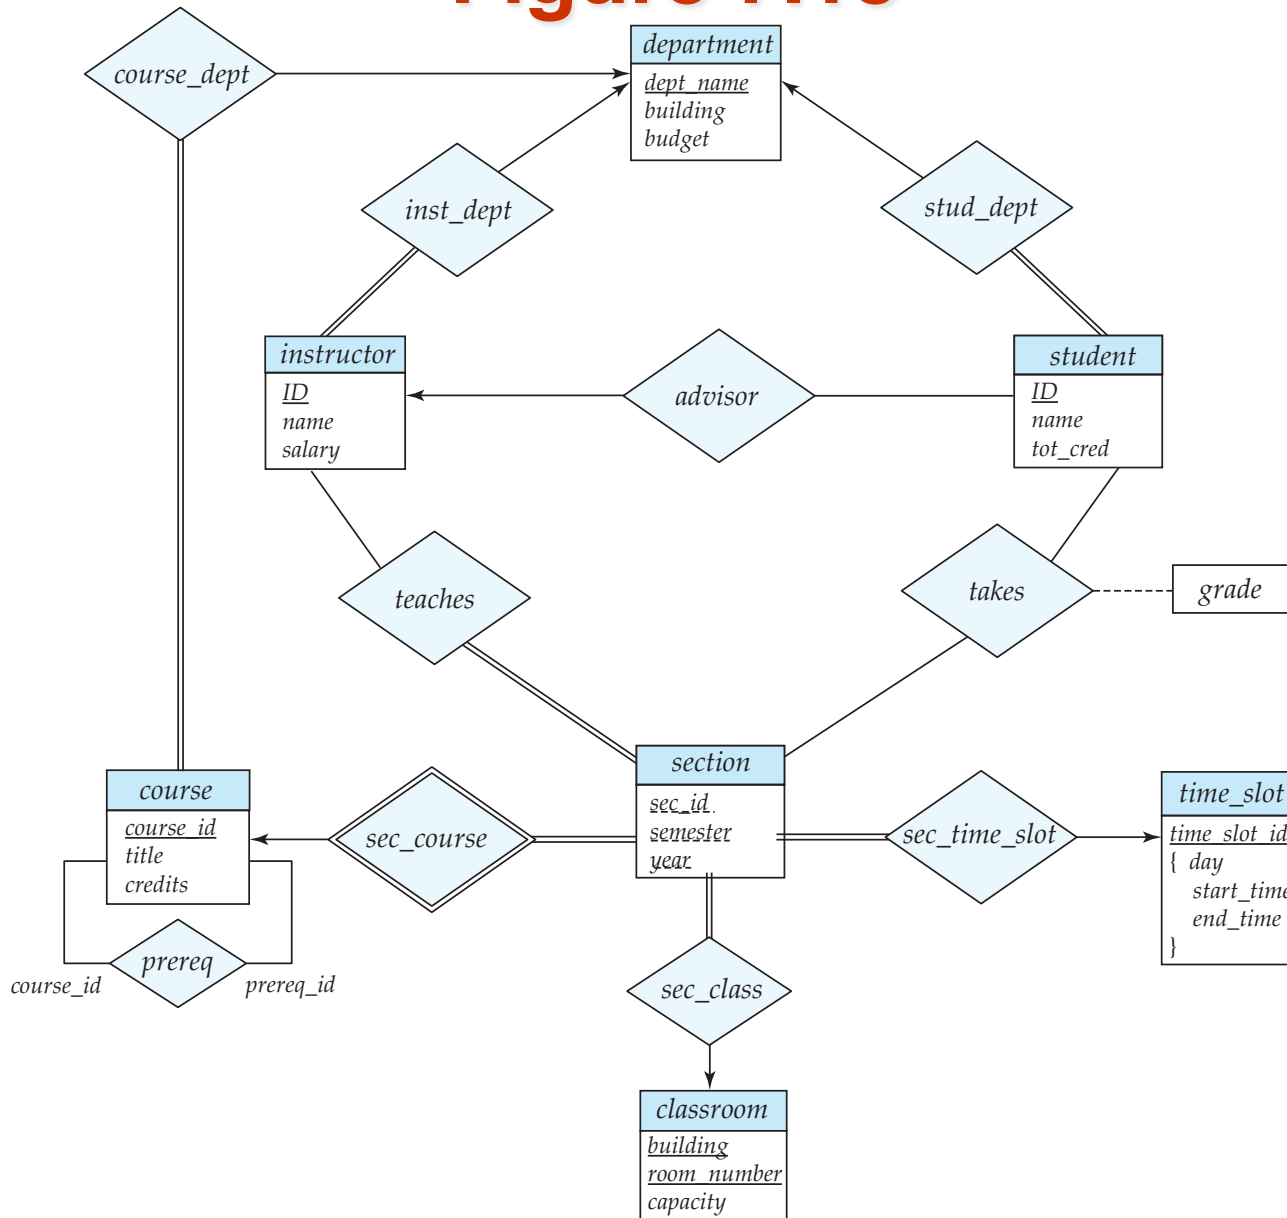
Figure 7.17

(a)

(b)

Figure 7.18

Figure 7.19

(a)

(b)

Figure 7.20

instructor

student
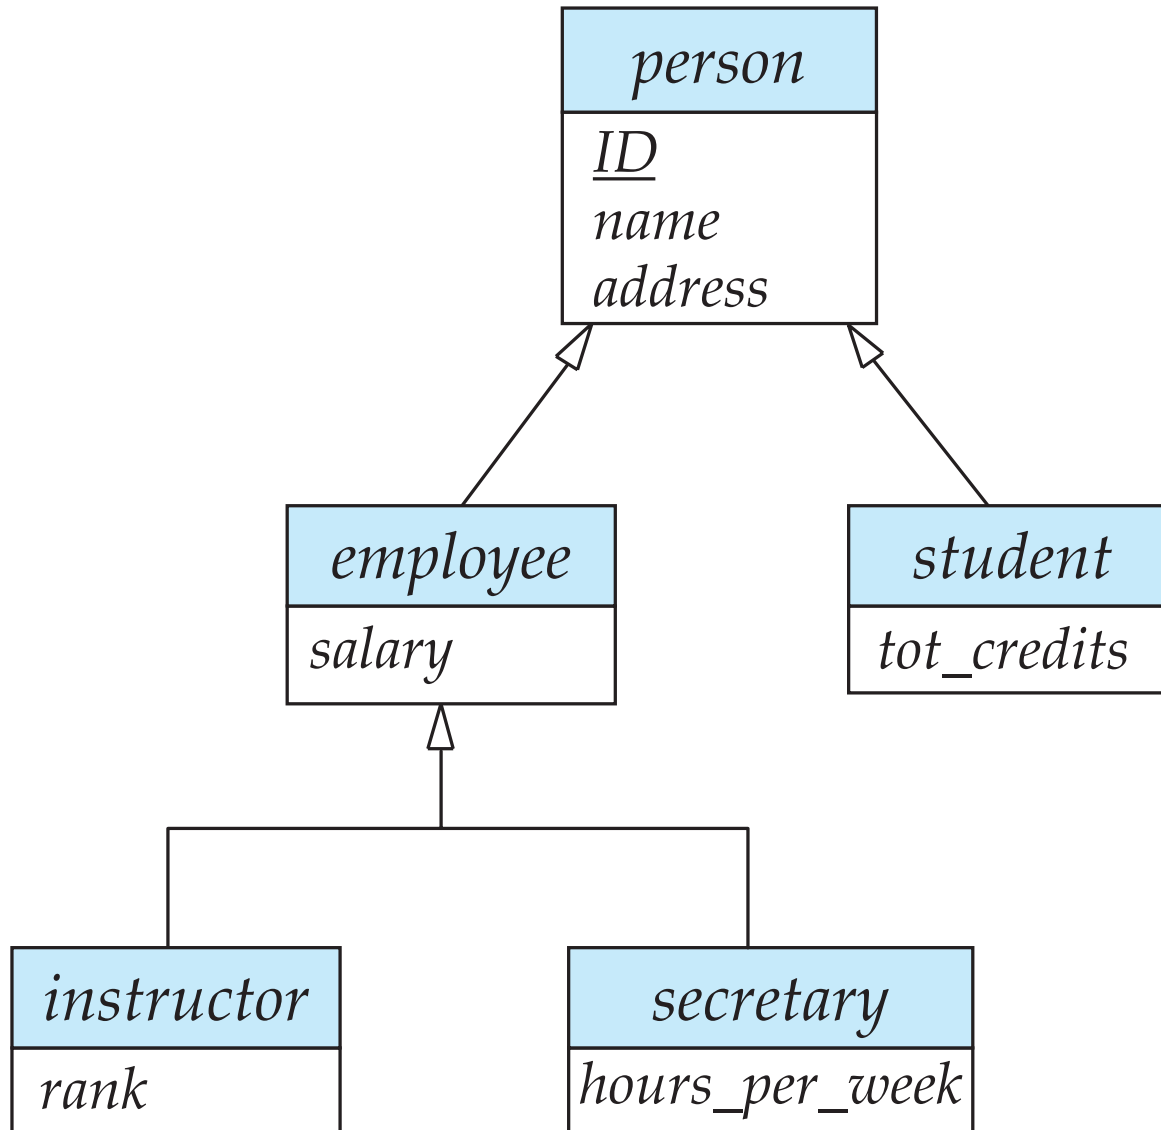
Figure 7.21

Figure 7.22

Figure 7.23

Figure 7.24

E	entity set

E	
A1	
A2	attributes:
A2.1	simple (A1),
A2.2	composite (A2) and
{A3}	multivalued (A3)
A4()	derived (A4)

R — relationship set

R (double diamond) — identifying relationship set for weak entity set

E	
A1	primary key

R = E — total participation of entity set in relationship

E	
A1	discriminating attribute of weak entity set

R (with lines) — many-to-many relationship

R (with arrow) — many-to-one relationship

R (with two arrows) — one-to-one relationship

R —l..h— E — cardinality limits

R —role-name— E — role indicator

E1 → E2, E3 (ISA triangles) — ISA: generalization or specialization

E1 ⇒ E2, E3 (total) — total (disjoint) generalization

E1 → E2, E3 — disjoint generalization

Figure 7.25

entity set E with
simple attribute A1,
composite attribute A2,
multivalued attribute A3,
derived attribute A4,
and primary key A1

many-to-many
relationship

one-to-one
relationship

many-to-one
relationship

participation
in R: total (E1)
and partial (E2)

weak entity set

generalization

total
generalization

Figure 7.26

Figure 7.27

Figure 7.28

Figure 7.29

CS425 – Fall 2013
Boris Glavic
Chapter 8: Relational Database Design

modified from:

Database System Concepts, 6th Ed.

Chapter 8: Relational Database Design

- Features of Good Relational Design
- Atomic Domains and First Normal Form
- Decomposition Using Functional Dependencies
- Functional Dependency Theory
- Algorithms for Functional Dependencies
- Decomposition Using Multivalued Dependencies
- More Normal Form
- Database-Design Process
- Modeling Temporal Data

What is Good Design?
1) Easier: What is Bad Design?

modified from:

Database System Concepts, 6th Ed.

Combine Schemas?

- Suppose we combine *instructor* and *department* into *inst_dept*
  - *(No connection to relationship set inst_dept)*
- Result is possible repetition of information

ID	name	salary	dept_name	building	budget
22222	Einstein	95000	Physics	Watson	70000
12121	Wu	90000	Finance	Painter	120000
32343	El Said	60000	History	Painter	50000
45565	Katz	75000	Comp. Sci.	Taylor	100000
98345	Kim	80000	Elec. Eng.	Taylor	85000
76766	Crick	72000	Biology	Watson	90000
10101	Srinivasan	65000	Comp. Sci.	Taylor	100000
58583	Califieri	62000	History	Painter	50000
83821	Brandt	92000	Comp. Sci.	Taylor	100000
15151	Mozart	40000	Music	Packard	80000
33456	Gold	87000	Physics	Watson	70000
76543	Singh	80000	Finance	Painter	120000

Redundancy is Bad!

ID	name	salary	dept_name	building	budget
22222	Einstein	95000	Physics	Watson	70000
12121	Wu	90000	Finance	Painter	120000
32343	El Said	60000	History	Painter	50000
45565	Katz	75000	Comp. Sci.	Taylor	100000
98345	Kim	80000	Elec. Eng.	Taylor	85000
76766	Crick	72000	Biology	Watson	90000
10101	Srinivasan	65000	Comp. Sci.	Taylor	100000
58583	Califieri	62000	History	Painter	50000
83821	Brandt	92000	Comp. Sci.	Taylor	100000
15151	Mozart	40000	Music	Packard	80000
33456	Gold	87000	Physics	Watson	70000
76543	Singh	80000	Finance	Painter	120000

- Update Physics Department
  - multiple tuples to update
  - Efficiency + potential for errors
- Delete Physics Department
  - update multiple tuples
  - Efficiency + potential for errors
- Departments without instructor or instructors without departments
  - Need dummy department and dummy instructor
  - Makes aggregation harder and error prone.

A Combined Schema Without Repetition

- Combining is not always bad!

- Consider combining relations

  - *sec_class(sec_id, building, room_number)* and

  - *section(course_id, sec_id, semester, year)*

  into one relation

  - *section(course_id, sec_id, semester, year, building, room_number)*

- No repetition in this case

What About Smaller Schemas?

- Suppose we had started with *inst_dept*. How would we know to split up (**decompose**) it into *instructor* and *department*?

- Write a rule "if there were a schema (*dept_name, building, budget*), then *dept_name* would be a candidate key"

- Denote as a **functional dependency**:

  dept_name → *building, budget*

- In *inst_dept*, because *dept_name* is not a candidate key, the building and budget of a department may have to be repeated.

  - This indicates the need to decompose *inst_dept*

- Not all decompositions are good. Suppose we decompose *employee(ID, name, street, city, salary)* into

  employee1 (*ID*, *name*)

  employee2 (*name*, *street*, *city*, *salary*)

- The next slide shows how we lose information -- we cannot reconstruct the original *employee* relation -- and so, this is a **lossy decomposition**.

A Lossy Decomposition
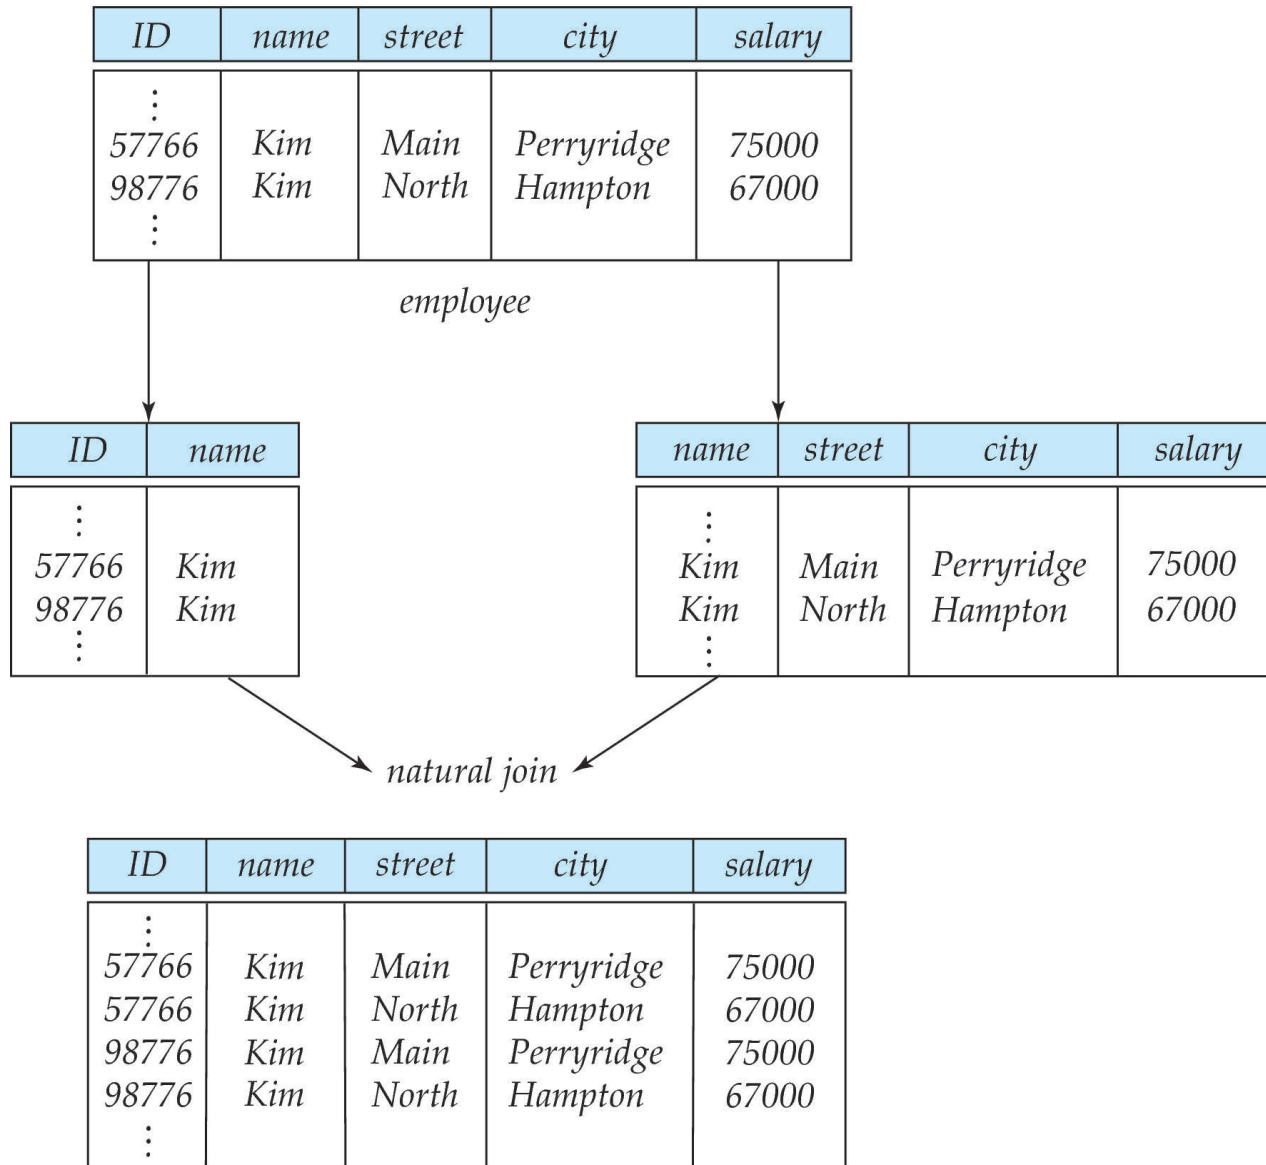
Example of Lossless-Join Decomposition

- **Lossless join decomposition**
- Decomposition of $R = (A, B, C)$

$$R_1 = (A, B) \qquad R_2 = (B, C)$$

A	B	C
α	1	A
β	2	B

r

A	B
α	1
β	2

$\prod_{A,B}(r)$

B	C
1	A
2	B

$\prod_{B,C}(r)$

$\prod_{A,B}(r) \bowtie \prod_{B,C}(r)$

A	B	C
α	1	A
β	2	B

Goals of Lossless-Join Decomposition

- Lossless-Join decomposition means splitting a table in a way so that we do not loose information

  - That means we should be able to reconstruct the original table from the decomposed table using joins

A	B	C
α	1	A
β	2	B

r

A	B
α	1
β	2

$\prod_{A,B}(r)$

B	C
1	A
2	B

$\prod_{B,C}(r)$

$$\prod_{A}(r) \bowtie \prod_{B}(r)$$

A	B	C
α	1	A
β	2	B

Goal — Devise a Theory for the Following

■ Decide whether a particular relation R is in "good" form.

■ In the case that a relation R is not in "good" form, decompose it into a set of relations $\{R_1, R_2, ..., R_n\}$ such that

● each relation is in good form

● the decomposition is a lossless-join decomposition

■ Our theory is based on:

● **1)** Models of dependency between attribute values

▸ **functional dependencies**

▸ multivalued dependencies

● **2)** Concept of **lossless decomposition**

● **3) Normal Forms** Based On

▸ Atomicity of values

▸ Avoidance of redundancy

▸ Lossless decomposition

Modeling Dependencies between Attribute Values:
Functional Depedencies
Multivalued Depedencies

modified from:

Database System Concepts, 6th Ed.

Functional Dependencies

■ Constraints on the set of legal instances for a relation schema.

■ Require that the value for a certain set of attributes determines uniquely the value for another set of attributes.

■ A functional dependency is a generalization of the notion of a *key.*

   ● *Thus, every key is a functional dependency*

Functional Dependencies (Cont.)

- Let R be a relation schema

$$\alpha \subseteq R \ \text{and} \ \beta \subseteq R$$

- The **functional dependency**

$$\alpha \rightarrow \beta$$

holds on R if and only if for any legal relations $r(R)$, whenever any two tuples t_1 and t_2 of r agree on the attributes α, they also agree on the attributes β. That is,

$$t_1[\alpha] = t_2[\alpha] \ \Rightarrow \ t_1[\beta] = t_2[\beta]$$

- Example: Consider $r(A,B)$ with the following instance of r.

1	4
1	5
3	7

- On this instance, $A \rightarrow B$ does **NOT** hold, but $B \rightarrow A$ does hold.

Functional Dependencies (Cont.)

- Let R be a relation schema

$$\alpha \subseteq R \;\; and \;\; \beta \subseteq R$$

- The **functional dependency**

$$\alpha \to \beta$$

  holds on R if and only if for any legal relations $r(R)$, whenever any two tuples t_1 and t_2 of r agree on the attributes α, they also agree on the attributes β. That is,

$$t_1[\alpha] = t_2[\alpha] \;\; \Rightarrow \;\; t_1[\beta] = t_2[\beta]$$

- Example: Consider $r(A,B)$ with the following instance of r.

1	4
1	5
3	7

A = 1 and B = 4
A = 1 and B = 5

- On this instance, $A \to B$ does **NOT** hold, but $B \to A$ does hold.

Functional Dependencies (Cont.)

- *K* is a superkey for relation schema *R* if and only if $K \rightarrow R$

- *K* is a candidate key for *R* if and only if
    - $K \rightarrow R$, and
    - for no $\alpha \subset K, \alpha \rightarrow R$

- Functional dependencies allow us to express constraints that cannot be expressed using superkeys. Consider the schema:

    inst_dept (*ID, name, salary, dept_name, building, budget*).

    We expect these functional dependencies to hold:

    $$dept\_name \rightarrow building$$

    and   ID → *building*

    but would not expect the following to hold:

    $$dept\_name \rightarrow salary$$

Use of Functional Dependencies

- We use functional dependencies to:
  - test relations to see if they are legal under a given set of functional dependencies.
    - If a relation *r* is legal under a set *F* of functional dependencies, we say that *r* **satisfies** *F*.
  - specify constraints on the set of legal relations
    - We say that *F* **holds on** *R* if all legal relations on *R* satisfy the set of functional dependencies *F*.
- Note: A specific instance of a relation schema may satisfy a functional dependency even if the functional dependency does not hold on all legal instances.
  - For example, a specific instance of *instructor* may, by chance, satisfy *name* → *ID*.

Functional Dependencies (Cont.)

■ *A* functional dependency is **trivial** if it is satisfied by all instances of a relation

- Example*:*
  - ▸ *ID, name → ID*
  - ▸ *name → name*
- In general, $\alpha \rightarrow \beta$ is trivial if $\beta \subseteq \alpha$

Closure of a Set of Functional Dependencies

- Given a set *F* of functional dependencies, there are certain other functional dependencies that are logically implied by *F*.

  - For example: If $A \rightarrow B$ and $B \rightarrow C$, then we can infer that $A \rightarrow C$

- The set of **all** functional dependencies logically implied by *F* is the **closure** of *F*.

- We denote the *closure* of *F* by $\textbf{F}^{+}$.

- F^{+} is a superset of *F*.

Functional-Dependency Theory

■ We now consider the formal theory that tells us which functional dependencies are implied logically by a given set of functional dependencies.

■ How do we get the initial set of FDs?

   ● Semantics of the domain we are modelling

   ● Has to be provided by a human (the designer)

■ Example:

   ● Relation Citizen(SSN, FirstName, LastName, Address)

   ● We know that SSN is unique and a person has a a unique SSN

   ● Thus, SSN → FirstName, LastName

Closure of a Set of Functional Dependencies

- We can find F^+, the closure of F, by repeatedly applying **Armstrong's Axioms:**

    - if $\beta \subseteq \alpha$, then $\alpha \rightarrow \beta$ (**reflexivity**)
    - if $\alpha \rightarrow \beta$, then $\gamma\,\alpha \rightarrow \gamma\,\beta$ (**augmentation**)
    - if $\alpha \rightarrow \beta$, and $\beta \rightarrow \gamma$, then $\alpha \rightarrow \gamma$ (**transitivity**)

- These rules are

    - **sound** (generate only functional dependencies that actually hold), and
    - **complete** (generate all functional dependencies that hold).

Example

- $R = (A, B, C, G, H, I)$
  $F = \{ \quad A \to B$
  $\qquad A \to C$
  $\qquad CG \to H$
  $\qquad CG \to I$
  $\qquad B \to H\}$

- some members of F^+

  - $A \to H$

    ‣ by transitivity from $A \to B$ and $B \to H$

  - $AG \to I$

    ‣ by augmenting $A \to C$ with G, to get $AG \to CG$
    and then transitivity with $CG \to I$

  - $CG \to HI$

    ‣ by augmenting $CG \to I$ to infer $CG \to CGI$,
    and augmenting of $CG \to H$ to infer $CGI \to HI$,
    and then transitivity

Prove Additional Implications

■ Prove or disprove the following rules from Amstrong's axioms

- *1) A → B, C implies A → B and A → C*

- *2) A → B and A → C implies A → B, C*

- *3) A, B → B, C implies A → C*

- 4) A → B and C → D implies A, C → B, D

Procedure for Computing F⁺

■ To compute the closure of a set of functional dependencies F:

$F^+ = F$

repeat

　　for each functional dependency f in F^+

　　　　apply reflexivity and augmentation rules on f

　　　　add the resulting functional dependencies to F^+

　　for each pair of functional dependencies f_1 and f_2 in F^+

　　　　if f_1 and f_2 can be combined using transitivity

　　　　　　then add the resulting functional dependency to F^+

until F^+ does not change any further

NOTE: We shall see an alternative more efficient procedure for this task later

Closure of Functional Dependencies (Cont.)

- Additional rules:

  - If $\alpha \rightarrow \beta$ holds and $\alpha \rightarrow \gamma$ holds, then $\alpha \rightarrow \beta\gamma$ holds **(union)**

  - If $\alpha \rightarrow \beta\gamma$ holds, then $\alpha \rightarrow \beta$ holds and $\alpha \rightarrow \gamma$ holds **(decomposition)**

  - If $\alpha \rightarrow \beta$ holds and $\gamma\beta \rightarrow \delta$ holds, then $\alpha\gamma \rightarrow \delta$ holds **(pseudotransitivity)**

  The above rules can be inferred from Armstrong's axioms.

Closure of Attribute Sets

- Given a set of attributes α, define the ***closure*** of α **under** *F* (denoted by α^+) as the set of attributes that are functionally determined by α under *F*

- Algorithm to compute α^+, the closure of α under *F*

> *result* := α;
> **while** (changes to *result*) **do**
>   **for each** $\beta \rightarrow \gamma$ **in** *F* **do**
>   **begin**
>   **if** $\beta \subseteq$ *result* **then** *result* := *result* $\cup$ γ
>   **end**

Example of Attribute Set Closure

- $R = (A, B, C, G, H, I)$
- $F = \{A \rightarrow B$
  $A \rightarrow C$
  $CG \rightarrow H$
  $CG \rightarrow I$
  $B \rightarrow H\}$

- $(AG)^+$

  1. $result = AG$

  2. $result = ABCG$ $(A \rightarrow C$ and $A \rightarrow B)$

  3. $result = ABCGH$ $(CG \rightarrow H$ and $CG \subseteq AGBC)$

  4. $result = ABCGHI$ $(CG \rightarrow I$ and $CG \subseteq AGBCH)$

- Is AG a candidate key?

  1. Is AG a super key?

     1. Does $AG \rightarrow R$? == Is $(AG)^+ \subseteq R$

  2. Is any subset of AG a superkey?

     1. Does $A \rightarrow R$? == Is $(A)^+ \subseteq R$

     2. Does $G \rightarrow R$? == Is $(G)^+ \subseteq R$

Uses of Attribute Closure

There are several uses of the attribute closure algorithm:

- Testing for superkey:
    - To test if α is a superkey, we compute α^+, and check if α^+ contains all attributes of R.

- Testing functional dependencies
    - To check if a functional dependency $\alpha \to \beta$ holds (or, in other words, is in F^+), just check if $\beta \subseteq \alpha^+$.
    - That is, we compute α^+ by using attribute closure, and then check if it contains β.
    - Is a simple and cheap test, and very useful

- Computing closure of F
    - For each $\gamma \subseteq R$, we find the closure γ^+, and for each $S \subseteq \gamma^+$, we output a functional dependency $\gamma \to S$.

O(n) Algorithm for Attribute Closure

- **Data Structures**

  - Enumerate the FDs and attributes

  - **int[] c**: an integer array with one element per FD that is initialized to the size of the LHS of the FD

  - **list<int>[] rhs**: an array of lists with one element per FD. The element stores the numeric ID of the attributes of the FDs RHS

  - **list<int>[] lhs**: an array of lists of integers, one element per attribute. The element for each attribute stores the numeric IDs of the FDs that have the attribute in its LHS

  - **set<int> aplus**: a set storing the attributes currently established to be implied by A

  - **stack<int> todo**: a stack of attributes to be processed next

O(n) Algorithm for Attribute Closure

- **Algorithm**
  - Initialize **c**, **rhs**, **lhs**, **aplus** to the emptyset, **todo** to **A**

```
while(!todo.isEmpty) {
    curA = todo.pop();
    aplus.add(curA);        // add curA to result
    for fd in lhs[curA]  { // update how many attribute found for LHS
        c[fd]--;            // found a LHS attr for fd
        if (c[fd] == 0) {
            remove(lhs[curA], fd); // avoid firing twice
            for newA in rhs[fd] { // add implied attributes
                if (!aplus[newA]) // if attribute is new add to todo
                    todo.push(newA);
                aplus.add(newA);
            }
        }
    }
}
```

Canonical Cover

- Sets of functional dependencies may have redundant dependencies that can be inferred from the others

    - For example: $A \rightarrow C$ is redundant in: $\{A \rightarrow B, \quad B \rightarrow C, A \rightarrow C\}$

    - Parts of a functional dependency may be redundant

        - E.g.: on RHS: $\{A \rightarrow B, \quad B \rightarrow C, \quad A \rightarrow CD\}$ can be simplified to

            $$\{A \rightarrow B, \quad B \rightarrow C, \quad A \rightarrow D\}$$

        - E.g.: on LHS: $\{A \rightarrow B, \quad B \rightarrow C, \quad AC \rightarrow D\}$ can be simplified to

            $$\{A \rightarrow B, \quad B \rightarrow C, \quad A \rightarrow D\}$$

- Intuitively, a **canonical cover** of F is a "minimal" set of functional dependencies equivalent to F, having no redundant dependencies or redundant parts of dependencies

Extraneous Attributes

- Consider a set F of functional dependencies and the functional dependency $\alpha \rightarrow \beta$ in F.

  - Attribute A is **extraneous** in α if $A \in \alpha$
    and F logically implies $(F - \{\alpha \rightarrow \beta\}) \cup \{(\alpha - A) \rightarrow \beta\}$.

  - Attribute A is **extraneous** in β if $A \in \beta$
    and the set of functional dependencies
    $(F - \{\alpha \rightarrow \beta\}) \cup \{\alpha \rightarrow (\beta - A)\}$ logically implies F.

- *Note:* implication in the opposite direction is trivial in each of the cases above, since a "stronger" functional dependency always implies a weaker one

- Example: Given $F = \{A \rightarrow C, AB \rightarrow C\}$

  - *B* is extraneous in $AB \rightarrow C$ because $\{A \rightarrow C, AB \rightarrow C\}$ logically implies $A \rightarrow C$ (I.e. the result of dropping B from $AB \rightarrow C$).

- Example: Given $F = \{A \rightarrow C, AB \rightarrow CD\}$

  - *C* is extraneous in $AB \rightarrow CD$ since $AB \rightarrow C$ can be inferred even after deleting C

Testing if an Attribute is Extraneous

- Consider a set F of functional dependencies and the functional dependency $\alpha \rightarrow \beta$ in F.

- To test if attribute $A \in \alpha$ is extraneous in α
    1. compute $(\{\alpha\} - A)^+$ using the dependencies in F
    2. check that $(\{\alpha\} - A)^+$ contains β; if it does, A is extraneous in α

- To test if attribute $A \in \beta$ is extraneous in β
    1. compute α^+ using only the dependencies in
        $$F' = (F - \{\alpha \rightarrow \beta\}) \cup \{\alpha \rightarrow (\beta - A)\},$$
    2. check that α^+ contains A; if it does, A is extraneous in β

Canonical Cover

- A **canonical cover** for F is a set of dependencies F_c such that
  - F logically implies all dependencies in F_c, and
  - F_c logically implies all dependencies in F, and
  - No functional dependency in F_c contains an extraneous attribute, and
  - Each left side of functional dependency in F_c is unique.

- To compute a canonical cover for F:
  repeat
     Use the union rule to replace any dependencies in F
           $\alpha_1 \rightarrow \beta_1$ and $\alpha_1 \rightarrow \beta_2$ with $\alpha_1 \rightarrow \beta_1 \beta_2$
     Find a functional dependency $\alpha \rightarrow \beta$ with an
           extraneous attribute either in α or in β
           /* Note: test for extraneous attributes done using F_c, not F*/
     If an extraneous attribute is found, delete it from $\alpha \rightarrow \beta$
  until F does not change

- Note: Union rule may become applicable after some extraneous attributes have been deleted, so it has to be re-applied

Computing a Canonical Cover

- $R = (A, B, C)$
  $F = \{A \rightarrow BC$
       $B \rightarrow C$
       $A \rightarrow B$
       $AB \rightarrow C\}$

- Combine $A \rightarrow BC$ and $A \rightarrow B$ into $A \rightarrow BC$
  - Set is now $\{A \rightarrow BC, B \rightarrow C, AB \rightarrow C\}$

- A is extraneous in $AB \rightarrow C$
  - Check if the result of deleting A from $AB \rightarrow C$ is implied by the other dependencies
    - Yes: in fact, $B \rightarrow C$ is already present!
  - Set is now $\{A \rightarrow BC, B \rightarrow C\}$

- C is extraneous in $A \rightarrow BC$
  - Check if $A \rightarrow C$ is logically implied by $A \rightarrow B$ and the other dependencies
    - Yes: using transitivity on $A \rightarrow B$ and $B \rightarrow C$.
      - Can use attribute closure of A in more complex cases

- The canonical cover is: $A \rightarrow B$
                                  $B \rightarrow C$

Lossless Join-Decomposition Dependency Preservation

modified from:

Database System Concepts, 6th Ed.

So Far

- **Theory of dependencies**

- **What is missing?**
  - When is a decomposition loss-less
    - ▸ Lossless-join decomposition
    - ▸ Dependencies on the input are preserved

- **What else is missing?**
  - Define what constitutes a good relation
    - ▸ Normal forms
  - How to check for a good relation
    - ▸ Test normal forms
  - How to achieve a good relation
    - ▸ Translate into normal form
    - ▸ Involves decomposition

Lossless-join Decomposition

- For the case of $R = (R_1, R_2)$, we require that for all possible relation instances r on schema R

$$r = \prod_{R1}(r) \bowtie \prod_{R2}(r)$$

- A decomposition of R into R_1 and R_2 is lossless join if at least one of the following dependencies is in F^+:

  - $R_1 \cap R_2 \to R_1$
  - $R_1 \cap R_2 \to R_2$

- The above functional dependencies are a **sufficient** condition for lossless join decomposition; the dependencies are a **necessary** condition only if all constraints are functional dependencies

Example

- $R = (A, B, C)$
  $F = \{A \to B, B \to C)$

  - Can be decomposed in two different ways

- $R_1 = (A, B), \quad R_2 = (B, C)$

  - Lossless-join decomposition:

  $$R_1 \cap R_2 = \{B\} \text{ and } B \to BC$$

  - Dependency preserving

- $R_1 = (A, B), \quad R_2 = (A, C)$

  - Lossless-join decomposition:

  $$R_1 \cap R_2 = \{A\} \text{ and } A \to AB$$

  - Not dependency preserving
    (cannot check $B \to C$ without computing $R_1 \bowtie R_2$)

Dependency Preservation

- Let F_i be the set of dependencies F^+ that include only attributes in R_i.

    - A decomposition is **dependency preserving**, if

        $$(F_1 \cup F_2 \cup \ldots \cup F_n)^+ = F^+$$

    - If it is not, then checking updates for violation of functional dependencies may require computing joins, which is expensive.

Testing for Dependency Preservation

- To check if a dependency $\alpha \to \beta$ is preserved in a decomposition of R into $R_1, R_2, \ldots, R_n$ we apply the following test (with attribute closure done with respect to F)

  - $result = \alpha$
    while (changes to $result$) do
      for each R_i in the decomposition
              $t = (result \cap R_i)^+ \cap R_i$
              $result = result \cup t$

  - If $result$ contains all attributes in β, then the functional dependency
    $\alpha \to \beta$ is preserved.

- We apply the test on all dependencies in F to check if a decomposition is dependency preserving

- This procedure (attribute closure) takes polynomial time, instead of the exponential time required to compute F^+ and $(F_1 \cup F_2 \cup \ldots \cup F_n)^+$

Example

- $R = (A, B, C)$
  $F = \{A \rightarrow B$
  $\quad\quad B \rightarrow C\}$
  Key = $\{A\}$

- Decomposition $R_1 = (A, B)$, $R_2 = (B, C)$
  - Lossless-join decomposition
  - Dependency preserving

Normal Forms

modified from:

Database System Concepts, 6th Ed.

So Far

- **Theory of dependencies**

- **Decompositions and ways to check whether they are "good"**
  - Lossless
  - Dependency preserving

- **What is missing?**
  - Define what constitutes a good relation
    - ▸ Normal forms
  - How to check for a good relation
    - ▸ Test normal forms
  - How to achieve a good relation
    - ▸ Translate into normal form
    - ▸ Involves decomposition

Goals of Normalization

- Let R be a relation scheme with a set F of functional dependencies.

- Decide whether a relation scheme R is in "good" form.

- In the case that a relation scheme R is not in "good" form, decompose it into a set of relation scheme $\{R_1, R_2, ..., R_n\}$ such that

  - each relation scheme is in good form

  - the decomposition is a lossless-join decomposition

  - Preferably, the decomposition should be dependency preserving.

First Normal Form

- A domain is **atomic** if its elements are considered to be indivisible units
  - Examples of non-atomic domains:
    - Set of names, composite attributes
    - Identification numbers like CS101 that can be broken up into parts
- A relational schema R is in **first normal form** if the domains of all attributes of R are atomic
- Non-atomic values complicate storage and encourage redundant (repeated) storage of data
  - Example: Set of accounts stored with each customer, and set of owners stored with each account
  - We assume all relations are in first normal form
  - (revisited in Chapter 22 of the textbook: Object Based Databases)

First Normal Form (Cont'd)

- Atomicity is actually a property of how the elements of the domain are used.

  - Example: Strings would normally be considered indivisible

  - Suppose that students are given roll numbers which are strings of the form *CS0012* or *EE1127*

  - If the first two characters are extracted to find the department, the domain of roll numbers is not atomic.

  - Doing so is a bad idea: leads to encoding of information in application program rather than in the database.

Second Normal Form

- A relation schema R in **1NF** is in **second normal form (2NF)** iff

  - No non-prime attribute depends on parts of a candidate key

  - An attribute is non-prime if it does not belong to any candidate key for R

Second Normal Form Example

- R(A,B,C,D)
  - A,B → C,D
  - A → C
  - B → D

- {A,B} is the only candidate key

- R is not in 2NF, because A->C where A is part of a candidate key and C is not part of a candidate key

- Interpretation **R**(A,B,C,D) is **Advisor**(InstrSSN, StudentCWID, InstrName, StudentName)
  - Indication that we are putting stuff together that does not belong together

Second Normal Form Interpretation

- Why is a dependency on parts of a candidate key bad?

  - That is why is a relation that is not in 2NF bad?

- 1) A dependency on part of a candidate key indicates potential for redudancy

  - **Advisor**(InstrSSN, StudentCWID, InstrName, StudentName)

  - StudentCWID → StudentName

  - If a student is advised by multiple instructors we record his name several times

- 2) A dependency on parts of a candidate key shows that some attributes are unrelated to other parts of a candidate key

  - That means the table should be split

2NF is What We Want?

- **Instructor**(Name, Salary, DepName, DepBudget) = I(A,B,C,D)
  - A → B,C,D
  - C → D
- {Name} is the only candidate key
- I is in 2NF
- However, as we have seen before I still has update redundancy that can cause update anomalies
  - We repeat the budget of a department if there is more than one instructor working for that department

Third Normal Form

■ A relation schema R is in **third normal form** **(3NF)** if for all:

$$\alpha \rightarrow \beta \text{ in } F^+$$

at least one of the following holds:

- $\alpha \rightarrow \beta$ is trivial (i.e., $\beta \in \alpha$)

- α is a superkey for R

- Each attribute A in $\beta - \alpha$ is contained in a candidate key for R.

  (**NOTE**: each attribute may be in a different candidate key)

Alternatively,

- Every attribute depends directly on a candidate key, i.e., for every attribute A there is a dependency $X \rightarrow A$, but no dependency $Y \rightarrow A$ where Y is not a candidate key

3NF Example

- **Instructor**(Name, Salary, DepName, DepBudget) = I(A,B,C,D)
  - $A \rightarrow B,C,D$
  - $C \rightarrow D$
- {Name} is the only candidate key
- I is in 2NF
- I is not in 3NF

Testing for 3NF

- Optimization: Need to check only FDs in F, need not check all FDs in F^+.

- Use attribute closure to check for each dependency $\alpha \rightarrow \beta$, if α is a superkey.

- If α is not a superkey, we have to verify if each attribute in β is contained in a candidate key of R
  - this test is rather more expensive, since it involve finding candidate keys
  - testing for 3NF has been shown to be NP-hard
  - Interestingly, decomposition into third normal form (described shortly) can be done in polynomial time

3NF Decomposition Algorithm

Let F_c be a canonical cover for F;
$i := 0$;
for each functional dependency $\alpha \to \beta$ in F_c **do**
 if none of the schemas R_j, $1 \le j \le i$ contains α β
      then begin
           $i := i + 1$;
           $R_i := \alpha$ β
      end
if none of the schemas R_j, $1 \le j \le i$ contains a candidate key for R
 then begin
           $i := i + 1$;
           $R_i :=$ any candidate key for R;
      end
/* Optionally, remove redundant relations */

 repeat
if any schema R_j is contained in another schema R_k
      then /* delete R_j */
           $R_j = R_i$;
           $i=i-1$;
return $(R_1, R_2, ..., R_i)$

3NF Decomposition Algorithm (Cont.)

- Above algorithm ensures:
  - each relation schema R_i is in 3NF
  - decomposition is dependency preserving and lossless-join
  - Proof of correctness is at end of this presentation (<u>click here</u>)

3NF Decomposition: An Example

- Relation schema:

  cust_banker_branch = (<u>customer_id, employee_id</u>, branch_name, type)

- The functional dependencies for this relation schema are:

  1. *customer_id, employee_id → branch_name, type*

  2. *employee_id → branch_name*

  3. *customer_id, branch_name → employee_id*

- We first compute a canonical cover

  - *branch_name* is extraneous in the r.h.s. of the 1st dependency

  - No other attribute is extraneous, so we get F_C =

    customer_id, employee_id → type
    employee_id → branch_name
    customer_id, branch_name → employee_id

3NF Decompsition Example (Cont.)

- The **for** loop generates following 3NF schema:

    (*customer_id, employee_id, type*)

    (*employee_id, branch_name*)

    (*customer_id, branch_name, employee_id)*

    - Observe that (*customer_id, employee_id, type*) contains a candidate key of the original schema, so no further relation schema needs be added

- At end of for loop, detect and delete schemas, such as (*employee_id, branch_name*), which are subsets of other schemas

    - result will not depend on the order in which FDs are considered

- The resultant simplified 3NF schema is:

    (*customer_id, employee_id, type*)

    (*customer_id, branch_name, employee_id)*

Another 3NF Example

- Relation *dept_advisor*:
  - ***dept_advisor*** (*s_ID, i_ID, dept_name)*
    $F = \{s\_ID, dept\_name \rightarrow i\_ID,$

    $i\_ID \rightarrow dept\_name\}$
  - Two candidate keys: *s_ID, dept_name,* and *i_ID, s_ID*
  - *R* is in 3NF
    - *s_ID, dept_name → i_ID   s_ID*
      - *dept_name* is a superkey
    - *i_ID → dept_name*
      - *dept_name* is contained in a candidate key

Redundancy in 3NF

- There is some redundancy in this schema **dept_advisor** (*s_ID, i_ID, dept_name*)

- Example of problems due to redundancy in 3NF

  - $R = (J, K, L)$
    $F = \{JK \rightarrow L, L \rightarrow K\}$

J	L	K
j_1	l_1	k_1
j_2	l_1	k_1
j_3	l_1	k_1
null	l_2	k_2

- repetition of information (e.g., the relationship l_1, k_1)

  - (*i_ID, dept_name*)

- need to use null values (e.g., to represent the relationship l_2, k_2 where there is no corresponding value for *J*).

  - (*i_ID, dept_nameI*) if there is no separate relation mapping instructors to departments

Boyce-Codd Normal Form

A relation schema R is in BCNF with respect to a set F of functional dependencies if for all functional dependencies in F^+ of the form

$$\alpha \to \beta$$

where $\alpha \subseteq R$ and $\beta \subseteq R$, at least one of the following holds:

- $\alpha \to \beta$ is trivial (i.e., $\beta \subseteq \alpha$)
- α is a superkey for R

Example schema *not* in BCNF:

   instr_dept (*ID, name, salary, dept_name, building, budget*)

because *dept_name$\to$ building, budget*
holds on *instr_dept,* but *dept_name* is not a superkey

BCNF and Dependency Preservation

■ If a relation is in BCNF it is in 3NF

■ Constraints, including functional dependencies, are costly to check in practice unless they pertain to only one relation

■ Because it is **not always** possible to achieve **both BCNF and dependency preservation**, we usually consider normally *third normal form.*

Testing for BCNF

- To check if a non-trivial dependency $\alpha \rightarrow \beta$ causes a violation of BCNF
    1. compute α^+ (the attribute closure of α), and
    2. verify that it includes all attributes of R, that is, it is a superkey of R.
- **Simplified test**: To check if a relation schema R is in BCNF, it suffices to check only the dependencies in the given set F for violation of BCNF, rather than checking all dependencies in F^+.
    - If none of the dependencies in F causes a violation of BCNF, then none of the dependencies in F^+ will cause a violation of BCNF either.
- However, **simplified test using only F is incorrect when testing a relation in a decomposition of R**
    - Consider $R = (A, B, C, D, E)$, with $F = \{ A \rightarrow B, BC \rightarrow D\}$
        - Decompose R into $R_1 = (A,B)$ and $R_2 = (A,C,D, E)$
        - Neither of the dependencies in F contain only attributes from (A,C,D,E) so we might be mislead into thinking R_2 satisfies BCNF.
        - In fact, dependency $AC \rightarrow D$ in F^+ shows R_2 is not in BCNF.

Testing Decomposition for BCNF

- To check if a relation R_i in a decomposition of R is in BCNF,

  - Either test R_i for BCNF with respect to the **restriction** of F to R_i (that is, all FDs in F^+ that contain only attributes from R_i)

  - or use the original set of dependencies F that hold on R, but with the following test:

    - for every set of attributes $\alpha \subseteq R_i$, check that α^+ (the attribute closure of α) either includes no attribute of R_i- α, or includes all attributes of R_i.

    - If the condition is violated by some $\alpha \rightarrow \beta$ in F, the dependency
      $$\alpha \rightarrow (\alpha^+ - \alpha) \cap R_i$$
      can be shown to hold on R_i, and R_i violates BCNF.

    - We use above dependency to decompose R_i

Decomposing a Schema into BCNF

- Suppose we have a schema R and a non-trivial dependency $\alpha \rightarrow \beta$ causes a violation of BCNF.

  We decompose R into:

  - $(\alpha \cup \beta)$
  - $(R - (\beta - \alpha))$

- In our example,

  - $\alpha = dept\_name$
  - $\beta = building, budget$

  and *inst_dept* is replaced by

  - $(\alpha \cup \beta) = (dept\_name, building, budget)$
  - $(R - (\beta - \alpha)) = (ID, name, salary, dept\_name)$

BCNF Decomposition Algorithm

$result := \{R\}$;
$done :=$ false;
compute F^+;
while (not *done)* **do**
  if (there is a schema R_i in *result* that is not in BCNF)
    then begin
        let $\alpha \rightarrow \beta$ be a nontrivial functional dependency that
          holds on R_i such that $\alpha \rightarrow R_i$ is not in F^+,
           and $\alpha \cap \beta = \varnothing$;
         $result := (result - R_i) \cup (R_i - \beta) \cup (\alpha, \beta)$;
        end
    else *done :=* **true;**

Note: each R_i is in BCNF, and decomposition is lossless-join.

Example of BCNF Decomposition

- $R = (A, B, C)$
  $F = \{A \rightarrow B$
         $B \rightarrow C\}$
  Key = $\{A\}$

- R is not in BCNF ($B \rightarrow C$ but B is not superkey)

- Decomposition

  - $R_1 = (B, C)$

  - $R_2 = (A, B)$

Example of BCNF Decomposition

- *class* (*course_id*, *title*, *dept_name*, *credits*, *sec_id*, *semester*, *year*, *building*, *room_number*, *capacity*, *time_slot_id*)

- Functional dependencies:
  - *course_id*→ *title*, *dept_name*, *credits*
  - *building*, *room_number*→*capacity*
  - *course_id*, *sec_id*, *semester*, *year*→*building*, *room_number*, *time_slot_id*

- A candidate key {*course_id*, *sec_id*, *semester*, *year*}.

- BCNF Decomposition:
  - *course_id*→ *title*, *dept_name*, *credits* holds
    - but *course_id* is not a superkey.
  - We replace *class* by:
    - *course*(*course_id*, *title*, *dept_name*, *credits*)
    - *class-1* (*course_id*, *sec_id*, *semester*, *year*, *building*, *room_number*, *capacity*, *time_slot_id*)

BCNF Decomposition (Cont.)

- *course* is in BCNF

  - How do we know this?

- *building*, *room_number* → *capacity* holds on *class-1*

  - but {*building*, *room_number*} is not a superkey for *class-1*.

  - We replace *class-1* by:

    - *classroom* (*building*, *room_number*, *capacity*)

    - *section* (*course_id, sec_id, semester, year, building, room_number, time_slot_id*)

- *classroom* and *section* are in BCNF.

BCNF and Dependency Preservation

It is not always possible to get a BCNF decomposition that is dependency preserving

- $R = (J, K, L)$
  $F = \{JK \rightarrow L$
  $\qquad L \rightarrow K\}$
  Two candidate keys = JK and JL

- R is not in BCNF

- Any decomposition of R will fail to preserve

$$JK \rightarrow L$$

This implies that testing for $JK \rightarrow L$ requires a join

How good is BCNF?

- There are database schemas in BCNF that do not seem to be sufficiently normalized

- Consider a relation

    inst_info (ID, child_name, phone)

  - where an instructor may have more than one phone and can have multiple children

ID	child_name	phone
99999 99999 99999 99999	David David William Willian	512-555-1234 512-555-4321 512-555-1234 512-555-4321

inst_info

How good is BCNF? (Cont.)

- There are no non-trivial functional dependencies and therefore the relation is in BCNF

- Insertion anomalies – i.e., if we add a phone 981-992-3443 to 99999, we need to add two tuples

    (99999, David,   981-992-3443)
    (99999, William, 981-992-3443)

How good is BCNF? (Cont.)

■ Therefore, it is better to decompose *inst_info* into:

inst_child

ID	child_name
99999	David
99999	David
99999	William
99999	Willian

inst_phone

ID	phone
99999	512-555-1234
99999	512-555-4321
99999	512-555-1234
99999	512-555-4321

This suggests the need for higher normal forms, such as Fourth Normal Form (4NF), which we shall see later.

Comparison of BCNF and 3NF

- It is always possible to decompose a relation into a set of relations that are in 3NF such that:
  - the decomposition is lossless
  - the dependencies are preserved

- It is always possible to decompose a relation into a set of relations that are in BCNF such that:
  - the decomposition is lossless
  - it may not be possible to preserve dependencies.

Summary Normal Forms

- BCNF -> 3NF -> 2NF -> 1NF

- **1NF**
    - atomic attributes

- **2NF**
    - no non-trivial dependencies of non-prime attributes on parts of the key

- **3NF**
    - no transitive non-trivial dependencies on the key

- **BCNF**
    - only non-trivial dependencies on a superkey

Design Goals Revisited

- Goal for a relational database design is:
  - BCNF.
  - Lossless join.
  - Dependency preservation.

- If we cannot achieve this, we accept one of
  - Lack of dependency preservation
  - Redundancy due to use of 3NF

- Interestingly, SQL does not provide a direct way of specifying functional dependencies other than superkeys.

  Can specify FDs using assertions, but they are expensive to test, (and currently not supported by any of the widely used databases!)

- Even if we had a dependency preserving decomposition, using SQL we would not be able to efficiently test a functional dependency whose left hand side is not a key.

Multivalued Dependencies and 4NF, 5NF

modified from:

Database System Concepts, 6th Ed.

Multivalued Dependencies

- Suppose we record names of children, and phone numbers for instructors:

  - *inst_child*(*ID*, *child_name*)

  - *inst_phone*(*ID*, *phone_number*)

- If we were to combine these schemas to get

  - *inst_info*(*ID*, *child_name*, *phone_number*)

  - Example data:
    (99999, David, 512-555-1234)
    (99999, David, 512-555-4321)
    (99999, William, 512-555-1234)
    (99999, William, 512-555-4321)

- This relation is in BCNF

  - Why?

Multivalued Dependencies (MVDs)

- Let R be a relation schema and let $\alpha \subseteq R$ and $\beta \subseteq R$. The **multivalued dependency**

$$\alpha \rightarrow\rightarrow \beta$$

holds on R if in any legal relation $r(R)$, for all pairs for tuples t_1 and t_2 in r such that $t_1[\alpha] = t_2[\alpha]$, there exist tuples t_3 and t_4 in r such that:

$$t_1[\alpha] = t_2[\alpha] = t_3[\alpha] = t_4[\alpha]$$
$$t_3[\beta] \quad\quad = \quad t_1[\beta]$$
$$t_3[R - \beta] = \quad t_2[R - \beta]$$
$$t_4[\beta] \quad\quad = \quad t_2[\beta]$$
$$t_4[R - \beta] = \quad t_1[R - \beta]$$

MVD (Cont.)

■ Tabular representation of $\alpha \twoheadrightarrow \beta$

	α	β	$R - \alpha - \beta$
t_1	$a_1 \ldots a_i$	$a_{i+1} \ldots a_j$	$a_{j+1} \ldots a_n$
t_2	$a_1 \ldots a_i$	$b_{i+1} \ldots b_j$	$b_{j+1} \ldots b_n$
t_3	$a_1 \ldots a_i$	$a_{i+1} \ldots a_j$	$b_{j+1} \ldots b_n$
t_4	$a_1 \ldots a_i$	$b_{i+1} \ldots b_j$	$a_{j+1} \ldots a_n$

Example

- Let R be a relation schema with a set of attributes that are partitioned into 3 nonempty subsets.

$$Y, Z, W$$

- We say that $Y \twoheadrightarrow Z$ (Y **multidetermines** Z) if and only if for all possible relations $r(R)$

$$< y_1, z_1, w_1 > \in r \text{ and } < y_1, z_2, w_2 > \in r$$

then

$$< y_1, z_1, w_2 > \in r \text{ and } < y_1, z_2, w_1 > \in r$$

- Note that since the behavior of Z and W are identical it follows that

$$Y \twoheadrightarrow Z \text{ if } Y \twoheadrightarrow W$$

Example (Cont.)

- In our example:

$$ID \twoheadrightarrow child\_name$$
$$ID \twoheadrightarrow phone\_number$$

- The above formal definition is supposed to formalize the notion that given a particular value of Y (*ID*) it has associated with it a set of values of Z *(child_name)* and a set of values of W *(phone_number)*, and these two sets are in some sense independent of each other.

- Note:

  - If $Y \rightarrow Z$ then $Y \twoheadrightarrow Z$

  - Indeed we have (in above notation) $Z_1 = Z_2$
    The claim follows.

Use of Multivalued Dependencies

- We use multivalued dependencies in two ways:

  1. To test relations to **determine** whether they are legal under a given set of functional and multivalued dependencies

  2. To specify **constraints** on the set of legal relations. We shall thus concern ourselves *only* with relations that satisfy a given set of functional and multivalued dependencies.

- If a relation *r* fails to satisfy a given multivalued dependency, we can construct a relations *r'* that does satisfy the multivalued dependency by adding tuples to *r.*

Theory of MVDs

■ From the definition of multivalued dependency, we can derive the following rule:

- If $\alpha \to \beta$, then $\alpha \to\to \beta$

That is, every functional dependency is also a multivalued dependency

■ The **closure** D^+ of *D* is the set of all functional and multivalued dependencies logically implied by *D*.

- We can compute D^+ from *D*, using the formal definitions of functional dependencies and multivalued dependencies.

- We can manage with such reasoning for very simple multivalued dependencies, which seem to be most common in practice

- For complex dependencies, it is better to reason about sets of dependencies using a system of inference rules (see Appendix C).

Fourth Normal Form

- A relation schema R is in **4NF** with respect to a set D of functional and multivalued dependencies if for all multivalued dependencies in D^+ of the form $\alpha \twoheadrightarrow \beta$, where $\alpha \subseteq R$ and $\beta \subseteq R$, at least one of the following hold:

  - $\alpha \twoheadrightarrow \beta$ is trivial (i.e., $\beta \subseteq \alpha$ or $\alpha \cup \beta = R$)

  - α is a superkey for schema R

- If a relation is in 4NF it is in BCNF

Restriction of Multivalued Dependencies

- The restriction of D to R_i is the set D_i consisting of
  - All functional dependencies in D^+ that include only attributes of R_i
  - All multivalued dependencies of the form

    $$\alpha \rightarrow\rightarrow (\beta \cap R_i)$$

    where $\alpha \subseteq R_i$ and $\alpha \rightarrow\rightarrow \beta$ is in D^+

4NF Decomposition Algorithm

result: = {R};
done := false;
compute D⁺;
Let D_i denote the restriction of D^+ to R_i

while (**not** *done*)
   if (there is a schema **R**$_i$ in *result* that is not in 4NF) **then**
     begin

       let $\alpha \rightarrow\rightarrow \beta$ be a nontrivial multivalued dependency that holds
        on R_i such that $\alpha \rightarrow R_i$ is not in D_i, and $\alpha \cap \beta = \phi$;
       result := (*result* - R_i) $\cup$ (R_i - β) $\cup$ (α, β);
     end
   else *done*:= true;

Note: each R_i is in 4NF, and decomposition is lossless-join

Example

- $R = (A, B, C, G, H, I)$

  $F = \{\ A \twoheadrightarrow B$

  $\qquad B \twoheadrightarrow HI$

  $\qquad CG \twoheadrightarrow H\ \}$

- R is not in 4NF since $A \twoheadrightarrow B$ and A is not a superkey for R

- Decomposition

  a) $R_1 = (A, B)$ (R_1 is in 4NF)

  b) $R_2 = (A, C, G, H, I)$ (R_2 is not in 4NF, decompose into R_3 and R_4)

  c) $R_3 = (C, G, H)$ (R_3 is in 4NF)

  d) $R_4 = (A, C, G, I)$ (R_4 is not in 4NF, decompose into R_5 and R_6)

    - $A \twoheadrightarrow B$ and $B \twoheadrightarrow HI \Rightarrow A \twoheadrightarrow HI$, (MVD transitivity), and

    - and hence $A \twoheadrightarrow I$ *(MVD restriction to R_4)*

  e) $R_5 = (A, I)$ (R_5 is in 4NF)

  f) $R_6 = (A, C, G)$ (R_6 is in 4NF)

Further Normal Forms

- **Join dependencies** generalize multivalued dependencies
  - lead to **project-join normal form (PJNF)** (also called **fifth normal form**)
- A class of even more general constraints, leads to a normal form called **domain-key normal form**.
- Problem with these generalized constraints: are hard to reason with, and no set of sound and complete set of inference rules exists.
- Hence rarely used

Final Thoughts on Design Process

modified from:

Database System Concepts, 6th Ed.

Overall Database Design Process

■ We have assumed schema R is given

- R could have been generated when converting an ER diagram to a set of tables.

- R could have been a single relation containing *all* attributes that are of interest (called **universal relation**).

- Normalization breaks R into smaller relations.

- R could have been the result of some ad hoc design of relations, which we then test/convert to normal form.

ER Model and Normalization

- When an ER diagram is carefully designed, identifying all entities correctly, the tables generated from the ER diagram should not need further normalization.

- However, in a real (imperfect) design, there can be functional dependencies from non-key attributes of an entity to other attributes of the entity

    - Example: an *employee* entity with attributes
    department_name and *building*,
    and a functional dependency
    department_name→ building

    - Good design would have made department an entity

- Functional dependencies from non-key attributes of a relationship set possible, but rare --- most relationships are binary

Denormalization for Performance

- May want to use non-normalized schema for performance

- For example, displaying *prereqs* along with *course_id,* and *title* requires join of *course* with *prereq*

- Alternative 1: Use denormalized relation containing attributes of *course* as well as *prereq* with all above attributes

  - faster lookup

  - extra space and extra execution time for updates

  - extra coding work for programmer and possibility of error in extra code

- Alternative 2: use a materialized view defined as

  course   *prereq*

  - Benefits and drawbacks same as above, except no extra coding work for programmer and avoids possible errors

Other Design Issues

- Some aspects of database design are not caught by normalization

- Examples of bad database design, to be avoided:

  Instead of *earnings* (*company_id, year, amount*), use

  - *earnings_2004, earnings_2005, earnings_2006*, etc., all on the schema (*company_id, earnings*).
    - Above are in BCNF, but make querying across years difficult and needs new table each year

  - *company_year* (*company_id, earnings_2004, earnings_2005, earnings_2006*)
    - Also in BCNF, but also makes querying across years difficult and requires new attribute each year.
    - Is an example of a **crosstab**, where values for one attribute become column names
    - Used in spreadsheets, and in data analysis tools

Recap

- Functional and Multi-valued Dependencies
  - Axioms
  - Closure
  - Minimal Cover
  - Attribute Closure
- Redundancy and lossless decomposition
- Normal-Forms
  - 1NF, 2NF, 3NF
  - BCNF
  - 4NF, 5NF

End of Chapter

modified from:

Database System Concepts, 6th Ed.

Proof of Correctness of 3NF Decomposition Algorithm

modified from:

Database System Concepts, 6th Ed.

Correctness of 3NF Decomposition Algorithm

■ 3NF decomposition algorithm is dependency preserving (since there is a relation for every FD in F_c)

■ Decomposition is lossless

   ● A candidate key (C) is in one of the relations R_i in decomposition

   ● Closure of candidate key under F_c must contain all attributes in R.

   ● Follow the steps of attribute closure algorithm to show there is only one tuple in the join result for each tuple in R_i

Claim: if a relation R_i is in the decomposition generated by the above algorithm, then R_i satisfies 3NF.

■ Let R_i be generated from the dependency $\alpha \rightarrow \beta$

■ Let $\gamma \rightarrow B$ be any non-trivial functional dependency on R_i. (We need only consider FDs whose right-hand side is a single attribute.)

■ Now, B can be in either β or α but not in both. Consider each case separately.

Correctness of 3NF Decomposition (Cont'd.)

- Case 1: If B in β:

  - If γ is a superkey, the 2nd condition of 3NF is satisfied

  - Otherwise α must contain some attribute not in γ

  - Since $\gamma \rightarrow B$ is in F^+ it must be derivable from F_c, by using attribute closure on γ.

  - Attribute closure not have used $\alpha \rightarrow \beta$. If it had been used, α must be contained in the attribute closure of γ, which is not possible, since we assumed γ is not a superkey.

  - Now, using $\alpha \rightarrow (\beta - \{B\})$ and $\gamma \rightarrow B$, we can derive $\alpha \rightarrow B$

    (since $\gamma \subseteq \alpha \beta$, and $B \notin \gamma$ since $\gamma \rightarrow B$ is non-trivial)

  - Then, B is extraneous in the right-hand side of $\alpha \rightarrow \beta$; which is not possible since $\alpha \rightarrow \beta$ is in F_c.

  - Thus, if B is in β then γ must be a superkey, and the second condition of 3NF must be satisfied.

Correctness of 3NF Decomposition (Cont'd.)

- Case 2: B is in α.

  - Since α is a candidate key, the third alternative in the definition of 3NF is trivially satisfied.

  - In fact, we cannot show that γ is a superkey.

  - This shows exactly why the third alternative is present in the definition of 3NF.

  Q.E.D.

Figure 8.02

ID	name	salary	dept_name	building	budget
22222	Einstein	95000	Physics	Watson	70000
12121	Wu	90000	Finance	Painter	120000
32343	El Said	60000	History	Painter	50000
45565	Katz	75000	Comp. Sci.	Taylor	100000
98345	Kim	80000	Elec. Eng.	Taylor	85000
76766	Crick	72000	Biology	Watson	90000
10101	Srinivasan	65000	Comp. Sci.	Taylor	100000
58583	Califieri	62000	History	Painter	50000
83821	Brandt	92000	Comp. Sci.	Taylor	100000
15151	Mozart	40000	Music	Packard	80000
33456	Gold	87000	Physics	Watson	70000
76543	Singh	80000	Finance	Painter	120000

Figure 8.03

Figure 8.04

A	B	C	D
a_1	b_1	c_1	d_1
a_1	b_2	c_1	d_2
a_2	b_2	c_2	d_2
a_2	b_3	c_2	d_3
a_3	b_3	c_2	d_4

Figure 8.05

building	room_number	capacity
Packard	101	500
Painter	514	10
Taylor	3128	70
Watson	100	30
Watson	120	50

Figure 8.06

Figure 8.14

dept_name	ID	street	city
Physics	22222	North	Rye
Physics	22222	Main	Manchester
Finance	12121	Lake	Horseneck

Figure 8.15

dept_name	ID	street	city
Physics	22222	North	Rye
Math	22222	Main	Manchester

Figure 8.17

A	B	C
a_1	b_1	c_1
a_1	b_1	c_2
a_2	b_1	c_1
a_2	b_1	c_3

Chapter 9: Transactions

modified from:

Database System Concepts, 6th Ed.

©Silberschatz, Korth and Sudarshan
See www.db-book.com for conditions on re-use

Chapter 9: Transactions

- Transaction Concept
- Transaction State
- Concurrent Executions
- Serializability
- Recoverability
- Implementation of Isolation
- Transaction Definition in SQL
- Testing for Serializability.

Transaction Concept

- A **transaction** is a *unit* of program execution that accesses and possibly updates various data items.

- E.g. transaction to transfer $50 from account A to account B:

    1. **read**(*A*)
    2. *A* := *A* − 50
    3. **write**(*A*)
    4. **read**(*B*)
    5. *B* := *B* + 50
    6. **write**(*B*)

- Two main issues to deal with:

    - **Recovery:** Failures of various kinds, such as hardware failures and system crashes

    - **Concurrent:** execution of multiple transactions

Example of Fund Transfer

- Transaction to transfer $50 from account A to account B:
    1. **read**(A)
    2. $A := A - 50$
    3. **write**(A)
    4. **read**(B)
    5. $B := B + 50$
    6. **write**($B)$

- **Atomicity requirement**
    - if the transaction fails after step 3 and before step 6, money will be "lost" leading to an inconsistent database state
        - ▸ Failure could be due to software or hardware
    - the system should ensure that updates of a partially executed transaction are not reflected in the database

- **Durability requirement** — once the user has been notified that the transaction has completed (i.e., the transfer of the $50 has taken place), the updates to the database by the transaction must persist even if there are software or hardware failures.

Example of Fund Transfer (Cont.)

- Transaction to transfer $50 from account A to account B:
  1. **read**(*A*)
  2. *A* := *A* – 50
  3. **write**(*A*)
  4. **read**(*B*)
  5. *B* := *B* + 50
  6. **write**(*B*)

- **Consistency requirement** in above example:
  - the sum of A and B is unchanged by the execution of the transaction

- In general, consistency requirements include
  - ▸ Explicitly specified integrity constraints such as primary keys and foreign keys
  - ▸ Implicit integrity constraints
    - – e.g. sum of balances of all accounts, minus sum of loan amounts must equal value of cash-in-hand
  - A transaction must see a consistent database.
  - During transaction execution the database may be temporarily inconsistent.
  - When the transaction completes successfully the database must be consistent
    - ▸ Erroneous transaction logic can lead to inconsistency

Example of Fund Transfer (Cont.)

- **Isolation requirement** — if between steps 3 and 6, another transaction T2 is allowed to access the partially updated database, it will see an inconsistent database (the sum $A + B$ will be less than it should be).

T1	T2
1. **read**(A)	
2. $A := A - 50$	
3. **write**(A)	
	read(A), read(B), print(A+B)
4. **read**(B)	
5. $B := B + 50$	
6. **write**(B	

- Isolation can be ensured trivially by running transactions **serially**

  - that is, one after the other.

- However, executing multiple transactions concurrently has significant benefits, as we will see later.

ACID Properties

A **transaction** is a unit of program execution that accesses and possibly updates various data items.To preserve the integrity of data the database system must ensure:

- **Atomicity.** Either all operations of the transaction are properly reflected in the database or none are.

- **Consistency.** Execution of a transaction in isolation preserves the consistency of the database.

- **Isolation.** Although multiple transactions may execute concurrently, each transaction must be unaware of other concurrently executing transactions. Intermediate transaction results must be hidden from other concurrently executed transactions.

  - That is, for every pair of transactions T_i and T_j, it appears to T_i that either T_j finished execution before T_i started, or T_j started execution after T_i finished.

- **Durability.** After a transaction completes successfully, the changes it has made to the database persist, even if there are system failures.

Transaction State

- **Active** – the initial state; the transaction stays in this state while it is executing

- **Partially committed** – after the final statement has been executed.

- **Failed** -- after the discovery that normal execution can no longer proceed.

- **Aborted** – after the transaction has been rolled back and the database restored to its state prior to the start of the transaction. Two options after it has been aborted:

  - restart the transaction
    - can be done only if no internal logical error
  - kill the transaction

- **Committed** – after successful completion.

Transaction Model

- Operations
  - Read(A) – read value of data item A
  - Write(A) – write a new value of data item A
  - Commit – commit changes of the transaction
  - Abort – Revert changes made by the transaction
- Data Items
  - Objects in the data base
  - Usually we consider tuples (rows) or disk pages

9.10

Concurrent Executions

- Multiple transactions are allowed to run concurrently in the system. Advantages are:

  - **increased processor and disk utilization**, leading to better transaction *throughput*

    ▸ E.g. one transaction can be using the CPU while another is reading from or writing to the disk

    ▸ In multi-processor systems each statement can use one or more CPUs

  - **reduced average response time** for transactions: short transactions need not wait behind long ones.

- **Concurrency control schemes** – mechanisms to achieve isolation

  - that is, to control the interaction among the concurrent transactions in order to prevent them from destroying the consistency of the database

Schedules

- **Schedule** – a sequences of instructions that specify the chronological order in which instructions of concurrent transactions are executed

  - a schedule for a set of transactions must consist of all instructions of those transactions

  - must preserve the order in which the instructions appear in each individual transaction.

- A transaction that successfully completes its execution will have a commit instructions as the last statement

  - by default transaction assumed to execute commit instruction as its last step

- A transaction that fails to successfully complete its execution will have an abort instruction as the last statement

Schedule 1

- Let T_1 transfer \$50 from A to B, and T_2 transfer 10% of the balance from A to B.

- A serial schedule in which T_1 is followed by T_2 :

T_1	T_2
read (A)	
$A := A - 50$	
write (A)	
read (B)	
$B := B + 50$	
write (B)	
commit	
	read (A)
	$temp := A * 0.1$
	$A := A - temp$
	write (A)
	read (B)
	$B := B + temp$
	write (B)
	commit

Schedule 2

- A serial schedule where T_2 is followed by T_1

T_1	T_2
	read (A)
	$temp := A * 0.1$
	$A := A - temp$
	write (A)
	read (B)
	$B := B + temp$
	write (B)
	commit
read (A)	
$A := A - 50$	
write (A)	
read (B)	
$B := B + 50$	
write (B)	
commit	

Schedule 3

- Let T_1 and T_2 be the transactions defined previously. The following schedule is not a serial schedule, but it is *equivalent* to Schedule 1.

T_1	T_2
read (A)	
$A := A - 50$	
write (A)	
	read (A)
	$temp := A * 0.1$
	$A := A - temp$
	write (A)
read (B)	
$B := B + 50$	
write (B)	
commit	
	read (B)
	$B := B + temp$
	write (B)
	commit

In Schedules 1, 2 and 3, the sum A + B is preserved.

Schedule 4

■ The following concurrent schedule does not preserve the value of ($A + B$).

T_1	T_2
read (A) $A := A - 50$	
	read (A) $temp := A * 0.1$ $A := A - temp$ write (A) read (B)
write (A) read (B) $B := B + 50$ write (B) commit	
	$B := B + temp$ write (B) commit

Serializability

- **Basic Assumption** – Each transaction preserves database consistency.

- Thus serial execution of a set of transactions preserves database consistency.

- A (possibly concurrent) schedule is serializable if it is equivalent to a serial schedule. Different forms of schedule equivalence give rise to the notions of:

  1. **conflict serializability**
  2. **view serializability**

Simplified view of transactions

- We ignore operations other than **read** and **write** instructions

- We assume that transactions may perform arbitrary computations on data in local buffers in between reads and writes.

- Our simplified schedules consist of only **read** and **write** instructions.

Conflicting Instructions

■ Instructions I_i and I_j of transactions T_i and T_j respectively, **conflict** if and only if there exists some item Q accessed by both I_i and I_j, and at least one of these instructions wrote Q.

1. $I_i = \textbf{read}(Q)$, $I_j = \textbf{read}(Q)$. I_i and I_j don't conflict.
2. $I_i = \textbf{read}(Q)$, $I_j = \textbf{write}(Q)$. They conflict.
3. $I_i = \textbf{write}(Q)$, $I_j = \textbf{read}(Q)$.   They conflict
4. $I_i = \textbf{write}(Q)$, $I_j = \textbf{write}(Q)$. They conflict

■ Intuitively, a conflict between I_i and I_j forces a (logical) temporal order between them.

● If I_i and I_j are consecutive in a schedule and they do not conflict, their results would remain the same even if they had been interchanged in the schedule.

Conflict Serializability

- If a schedule S can be transformed into a schedule S' by a series of swaps of non-conflicting instructions, we say that S and S' are **conflict equivalent**.

  - That is the order of each pair of conflicting operations in S and S` is the same

- We say that a schedule S is **conflict serializable** if it is conflict equivalent to a serial schedule

Conflict Serializability (Cont.)

■ Schedule 3 can be transformed into Schedule 6, a serial schedule where T_2 follows T_1, by series of swaps of non-conflicting instructions. Therefore Schedule 3 is conflict serializable.

T_1	T_2
read (A)	
write (A)	
	read (A)
	write (A)
read (B)	
write (B)	
	read (B)
	write (B)

Schedule 3

T_1	T_2
read (A)	
write (A)	
read (B)	
write (B)	
	read (A)
	write (A)
	read (B)
	write (B)

Schedule 6

Conflict Serializability (Cont.)

- Example of a schedule that is not conflict serializable:

T_3	T_4
read (Q)	
	write (Q)
write (Q)	

- We are unable to swap instructions in the above schedule to obtain either the serial schedule < T_3, T_4 >, or the serial schedule < T_4, T_3 >.

View Serializability

■ Let S and S' be two schedules with the same set of transactions. S and S' are **view equivalent** if the following three conditions are met, for each data item Q,

1. If in schedule S, transaction T_i reads the initial value of Q, then in schedule S' also transaction T_i must read the initial value of Q.

2. If in schedule S transaction T_i executes **read**(*Q*), and that value was produced by transaction T_j (if any), then in schedule S' also transaction T_i must read the value of Q that was produced by the same **write**(Q) operation of transaction T_j .

3. The transaction (if any) that performs the final **write**(*Q*) operation in schedule S must also perform the final **write**(*Q*) operation in schedule S'.

As can be seen, view equivalence is also based purely on **reads** and **writes** alone.

View Serializability (Cont.)

- A schedule *S* is **view serializable** if it is view equivalent to a serial schedule.

- Every conflict serializable schedule is also view serializable.

- Below is a schedule which is view-serializable but *not* conflict serializable.

T_{27}	T_{28}	T_{29}
read (Q)		
	write (Q)	
write (Q)		
		write (Q)

- What serial schedule is above equivalent to?

- Every view serializable schedule that is not conflict serializable has **blind writes.**

Other Notions of Serializability

- The schedule below produces same outcome as the serial schedule $< T_1, T_5 >$, yet is not conflict equivalent or view equivalent to it.

T_1	T_5
read (A) $A := A - 50$ write (A)	
	read (B) $B := B - 10$ write (B)
read (B) $B := B + 50$ write (B)	
	read (A) $A := A + 10$ write (A)

- Determining such equivalence requires analysis of operations other than read and write.

Testing for Serializability

- Consider some schedule of a set of transactions T_1, T_2, ..., T_n

- **Precedence graph** — a directed graph where the vertices are the transactions (names).

- We draw an arc from T_i to T_j if the two transaction conflict, and T_i accessed the data item on which the conflict arose earlier.

- We may label the arc by the item that was accessed.

- **Example 1**

Test for Conflict Serializability

- A schedule is conflict serializable if and only if its precedence graph is acyclic.

- Cycle-detection algorithms exist which take order n^2 time, where n is the number of vertices in the graph.

  - (Better algorithms take order $n + e$ where e is the number of edges.)

- If precedence graph is acyclic, the serializability order can be obtained by a *topological sorting* of the graph.

  - This is a linear order consistent with the partial order of the graph.

  - For example, a serializability order for Schedule A would be
    $$T_5 \rightarrow T_1 \rightarrow T_3 \rightarrow T_2 \rightarrow T_4$$
    - Are there others?

(a)

(b)

(c)

Test for View Serializability

- The precedence graph test for conflict serializability cannot be used directly to test for view serializability.

  - Extension to test for view serializability has cost exponential in the size of the precedence graph.

- The problem of checking if a schedule is view serializable falls in the class of *NP*-complete problems.

  - Thus existence of an efficient algorithm is *extremely* unlikely.

- However practical algorithms that just check some **sufficient conditions** for view serializability can still be used.

Recoverable Schedules

Need to address the effect of transaction failures on concurrently running transactions.

- **Recoverable schedule** — if a transaction T_j reads a data item previously written by a transaction T_i, then the commit operation of T_i appears before the commit operation of T_j.

- The following schedule (Schedule 11) is not recoverable if T_9 commits immediately after the read

T_8	T_9
read (A)	
write (A)	
	read (A)
	commit
read (B)	

- If T_8 should abort, T_9 would have read (and possibly shown to the user) an inconsistent database state. Hence, database must ensure that schedules are recoverable.

Cascading Rollbacks

■ **Cascading rollback** – a single transaction failure leads to a series of transaction rollbacks. Consider the following schedule where none of the transactions has yet committed (so the schedule is recoverable)

T_{10}	T_{11}	T_{12}
read (A) read (B) write (A)		
	read (A) write (A)	
		read (A)
abort		

If T_{10} fails, T_{11} and T_{12} must also be rolled back.

■ Can lead to the undoing of a significant amount of work

Cascadeless Schedules

- **Cascadeless schedules** — cascading rollbacks cannot occur; for each pair of transactions T_i and T_j such that T_j reads a data item previously written by T_i, the commit operation of T_i appears before the read operation of T_j.

- Every cascadeless schedule is also recoverable

- It is desirable to restrict the schedules to those that are cascadeless

Concurrency Control

- A database must provide a mechanism that will ensure that all possible schedules are
    - either conflict or view serializable, and
    - are recoverable and preferably cascadeless
- A policy in which only one transaction can execute at a time generates serial schedules, but provides a poor degree of concurrency
    - Are serial schedules recoverable/cascadeless?
- Testing a schedule for serializability *after* it has executed is a little too late!
- **Goal** – to develop concurrency control protocols that will assure serializability.

Concurrency Control (Cont.)

■ Schedules must be conflict or view serializable, and recoverable, for the sake of database consistency, and preferably cascadeless.

■ A policy in which only one transaction can execute at a time generates serial schedules, but provides a poor degree of concurrency.

■ Concurrency-control schemes tradeoff between the amount of concurrency they allow and the amount of overhead that they incur.

■ Some schemes allow only conflict-serializable schedules to be generated, while others allow view-serializable schedules that are not conflict-serializable.

Concurrency Control vs. Serializability Tests

- Concurrency-control protocols allow concurrent schedules, but ensure that the schedules are conflict/view serializable, and are recoverable and cascadeless .

- Concurrency control protocols generally do not examine the precedence graph as it is being created
  - Instead a protocol imposes a discipline that avoids nonseralizable schedules.
  - We study such protocols in Chapter 10.

- Different concurrency control protocols provide different tradeoffs between the amount of concurrency they allow and the amount of overhead that they incur.

- Tests for serializability help us understand why a concurrency control protocol is correct.

Weak Levels of Consistency

- Some applications are willing to live with weak levels of consistency, allowing schedules that are not serializable

  - E.g. a read-only transaction that wants to get an approximate total balance of all accounts

  - E.g. database statistics computed for query optimization can be approximate (why?)

  - Such transactions need not be serializable with respect to other transactions

- Tradeoff accuracy for performance

Levels of Consistency in SQL-92

■ **Serializable** — default

■ **Repeatable read** — only committed records to be read, repeated reads of same record must return same value. However, a transaction may not be serializable – it may find some records inserted by a transaction but not find others.

■ **Read committed** — only committed records can be read, but successive reads of record may return different (but committed) values.

■ **Read uncommitted** — even uncommitted records may be read.

■ Lower degrees of consistency useful for gathering approximate information about the database

■ Warning: some database systems do not ensure serializable schedules by default

● E.g. Oracle and PostgreSQL by default support a level of consistency called snapshot isolation (not part of the SQL standard)

Transaction Definition in SQL

- Data manipulation language must include a construct for specifying the set of actions that comprise a transaction.

- In SQL, a transaction begins implicitly.

- A transaction in SQL ends by:
  - **Commit work** commits current transaction and begins a new one.
  - **Rollback work** causes current transaction to abort.

- In almost all database systems, by default, every SQL statement also commits implicitly if it executes successfully
  - Implicit commit can be turned off by a database directive
    - E.g. in JDBC,   connection.setAutoCommit(false);

End of Chapter 10

modified from:

Database System Concepts, 6th Ed.

Recap

- Transactions
  - ACID – Properties
- Schedules
  - Serial
  - Equivalence
    - ▸ Conflict-equivalent
    - ▸ View-equivalent
  - Serializability
    - ▸ = Equivalent to a serial schedule
  - Recoverable
  - Cascading Aborts
- Transactions in SQL

Figure 14.01

Figure 14.02

T_1	T_2
read (A)	
$A := A - 50$	
write (A)	
read (B)	
$B := B + 50$	
write (B)	
commit	
	read (A)
	$temp := A * 0.1$
	$A := A - temp$
	write (A)
	read (B)
	$B := B + temp$
	write (B)
	commit

Figure 14.03

T_1	T_2
	read (A)
	$temp := A * 0.1$
	$A := A - temp$
	write (A)
	read (B)
	$B := B + temp$
	write (B)
	commit
read (A)	
$A := A - 50$	
write (A)	
read (B)	
$B := B + 50$	
write (B)	
commit	

Figure 14.04

T_1	T_2
read (A) $A := A - 50$ write (A)	
	read (A) $temp := A * 0.1$ $A := A - temp$ write (A)
read (B) $B := B + 50$ write (B) commit	
	read (B) $B := B + temp$ write (B) commit

Figure 14.05

T_1	T_2
read (A) $A := A - 50$	
	read (A) $temp := A * 0.1$ $A := A - temp$ write (A) read (B)
write (A) read (B) $B := B + 50$ write (B) commit	
	$B := B + temp$ write (B) commit

Figure 14.06

T_1	T_2
read (A) write (A)	
	read (A) write (A)
read (B) write (B)	
	read (B) write (B)

Figure 14.07

T_1	T_2
read (A)	
write (A)	
	read (A)
read (B)	
	write (A)
write (B)	
	read (B)
	write (B)

Figure 14.08

T_1	T_2
read (A)	
write (A)	
read (B)	
write (B)	
	read (A)
	write (A)
	read (B)
	write (B)

Figure 14.09

T_3	T_4
read (Q)	
	write (Q)
write (Q)	

Figure 14.10

(a)
(b)

Figure 14.11

Figure 14.12

(a)

(b)

(c)

Figure 14.13

T_1	T_5
read (A) $A := A - 50$ write (A)	
	read (B) $B := B - 10$ write (B)
read (B) $B := B + 50$ write (B)	
	read (A) $A := A + 10$ write (A)

Figure 14.14

T_8	T_9
read (A)	
write (A)	
	read (A)
	commit
read (B)	

Figure 14.15

T_{10}	T_{11}	T_{12}
read (A) read (B) write (A)		
	read (A) write (A)	
		read (A)
abort		

Figure 14.16

Chapter 10 : Concurrency Control

modified from:

Database System Concepts, 6th Ed.

Chapter 10: Concurrency Control

- **Lock-Based Protocols**
- Timestamp-Based Protocols
- Validation-Based Protocols
- Multiple Granularity
- Multiversion Schemes
- Insert and Delete Operations
- Concurrency in Index Structures

Intuition of Lock-based Protocols

- Transactions have to acquire locks on data items before accessing them

- If a lock is hold by one transaction on a data item this restricts the ability of other transactions to acquire locks for that data item

- By locking a data item we want to ensure that no access to that data item is possible that would lead to non-serializable schedules

- The trick is to design a lock model and protocol that guarantees that

- Lock-based concurrency protocols are a form of **pessimistic concurrency control mechanism**

  - We avoid ever getting into a state that can lead to a non-serializable schedule

- Alternative concurrency control mechanism do not avoid conflicts, but determine later on (at commit time) whether committing a transaction would cause a non-serializable schedule to be generated

  - **Optimistic concurrency control mechanism**

Lock-Based Protocols

■ A lock is a mechanism to control concurrent access to a data item

■ Data items can be locked in two modes :

1. *exclusive (X) mode*. Data item can be both read as well as written. X-lock is requested using **lock-X** instruction.

2. *shared (S) mode*. Data item can only be read. S-lock is requested using **lock-S** instruction.

■ Lock requests are made to concurrency-control manager.

- Transaction do not access data items before having acquired a lock on that data item

- Transactions release their locks on a data item only after they have accessed a data item

Lock-Based Protocols (Cont.)

- **Lock-compatibility matrix**

	S	X
S	true	false
X	false	false

- A transaction may be granted a lock on an item if the requested lock is compatible with locks already held on the item by other transactions

- Any number of transactions can hold shared locks on an item,
  - but if any transaction holds an exclusive lock on the item no other transaction may hold any lock on the item.

- If a lock cannot be granted, the requesting transaction is made to wait till all incompatible locks held by other transactions have been released. The lock is then granted.

Lock-Based Protocols (Cont.)

- Example of a transaction performing locking:

  T_2: **lock-S**(A);

  read (A);

  unlock(A);

  lock-S(B);

  read (B);

  unlock(B);

  display$(A+B)$

- Locking as above is not sufficient to guarantee serializability — if A and B get updated in-between the read of A and B, the displayed sum would be wrong.

- A **locking protocol** is a set of rules followed by all transactions while requesting and releasing locks. Locking protocols restrict the set of possible schedules.

Pitfalls of Lock-Based Protocols

- Consider the partial schedule

T_3	T_4
lock-x (B)	
read (B)	
B := B - 50	
write (B)	
	lock-s (A)
	read (A)
	lock-s (B)
lock-x (A)	

- Neither T_3 nor T_4 can make progress — executing **lock-S**(B) causes T_4 to wait for T_3 to release its lock on B, while executing **lock-X**(A) causes T_3 to wait for T_4 to release its lock on A.

- Such a situation is called a **deadlock**.

  - To handle a deadlock one of T_3 or T_4 must be rolled back and its locks released.

Pitfalls of Lock-Based Protocols (Cont.)

- The potential for deadlock exists in most locking protocols. Deadlocks are a necessary evil.

- **Starvation** is also possible if the concurrency control manager is badly designed. For example:

  - A transaction may be waiting for an X-lock on an item, while a sequence of other transactions request and are granted an S-lock on the same item.

  - The same transaction is repeatedly rolled back due to deadlocks.

- Concurrency control managers can be designed to prevent starvation.

The Two-Phase Locking Protocol

- This is a protocol which ensures conflict-serializable schedules.

- Phase 1: Growing Phase
  - transaction may obtain locks
  - transaction may not release locks

- Phase 2: Shrinking Phase
  - transaction may release locks
  - transaction may not obtain locks

- The protocol assures serializability. It can be proved that the transactions can be serialized in the order of their **lock points** (i.e. the point where a transaction acquired its final lock).

The Two-Phase Locking Protocol (Cont.)

- Two-phase locking *does not* ensure freedom from deadlocks

- Cascading roll-back is possible under two-phase locking. To avoid this, follow a modified protocol called **strict two-phase locking**. Here a transaction must hold all its exclusive locks till it commits/aborts.

- **Rigorous two-phase locking** is even stricter: here *all* locks are held till commit/abort. In this protocol transactions can be serialized in the order in which they commit.

The Two-Phase Locking Protocol (Cont.)

■ There can be conflict serializable schedules that cannot be obtained if two-phase locking is used.

■ However, in the absence of extra information (e.g., ordering of access to data), two-phase locking is needed for conflict serializability in the following sense:

Given a transaction T_i that does not follow two-phase locking, we can find a transaction T_j that uses two-phase locking, and a schedule for T_i and T_j that is not conflict serializable.

Lock Conversions

- Two-phase locking with lock conversions:
  - First Phase:
    - can acquire a lock-S on item
    - can acquire a lock-X on item
    - can convert a lock-S to a lock-X (upgrade)
  - Second Phase:
    - can release a lock-S
    - can release a lock-X
    - can convert a lock-X to a lock-S (downgrade)
- This protocol assures serializability. But still relies on the programmer to insert the various locking instructions.

Automatic Acquisition of Locks

- A transaction T_i issues the standard read/write instruction, without explicit locking calls.

- The operation **read**(D) is processed as:

  if T_i has a lock on D

    then

      read(D)

    else begin

      if necessary wait until no other

        transaction has a **lock-X** on D

      grant T_i a **lock-S** on D;

      read(D)

    end

Automatic Acquisition of Locks (Cont.)

■ **write***(D)* is processed as:

if T_i has a **lock-X** on *D*

  then

    write(*D*)

  else begin

     if necessary wait until no other trans. has any lock on *D*,

     if T_i has a **lock-S** on *D*

       then

         upgrade lock on *D* to **lock-X**

      else

        grant T_i a **lock-X** on *D*

     write(*D*)

  end;

■ All locks are released after commit or abort

Implementation of Locking

- A **lock manager** can be implemented as a separate process to which transactions send lock and unlock requests

- The lock manager replies to a lock request by sending a lock grant messages (or a message asking the transaction to roll back, in case of a deadlock)

- The requesting transaction waits until its request is answered

- The lock manager maintains a data-structure called a **lock table** to record granted locks and pending requests

- The lock table is usually implemented as an in-memory hash table indexed on the name of the data item being locked

Lock Table

- Black rectangles indicate granted locks, white ones indicate waiting requests

- Lock table also records the type of lock granted or requested

- New request is added to the end of the queue of requests for the data item, and granted if it is compatible with all earlier locks

- Unlock requests result in the request being deleted, and later requests are checked to see if they can now be granted

- If transaction aborts, all waiting or granted requests of the transaction are deleted

  - lock manager may keep a list of locks held by each transaction, to implement this efficiently

granted

waiting

Graph-Based Protocols

■ Graph-based protocols are an alternative to two-phase locking

■ Impose a partial ordering → on the set **D** = {$d_1, d_2, ..., d_h$} of all data items.

  ● If $d_i → d_j$ then any transaction accessing both d_i and d_j must access d_i before accessing d_j.

  ● Implies that the set **D** may now be viewed as a directed acyclic graph, called a *database graph*.

■ The *tree-protocol* is a simple kind of graph protocol.

Tree Protocol

1. Only exclusive locks are allowed.

2. The first lock by T_i may be on any data item. Subsequently, a data Q can be locked by T_i only if the parent of Q is currently locked by T_i.

3. Data items may be unlocked at any time.

4. A data item that has been locked and unlocked by T_i cannot subsequently be relocked by T_i

Graph-Based Protocols (Cont.)

- The tree protocol ensures conflict serializability as well as freedom from deadlock.

- Unlocking may occur earlier in the tree-locking protocol than in the two-phase locking protocol.

    - shorter waiting times, and increase in concurrency

    - protocol is deadlock-free, no rollbacks are required

- Drawbacks

    - Protocol does not guarantee recoverability or cascade freedom

        ▸ Need to introduce commit dependencies to ensure recoverability

    - Transactions may have to lock data items that they do not access.

        ▸ increased locking overhead, and additional waiting time

        ▸ potential decrease in concurrency

- Schedules not possible under two-phase locking are possible under tree protocol, and vice versa.

Deadlock Handling

- Consider the following two transactions:

  T_1:   write (X) T_2:   write(Y)

          write(Y)   write(X)

- Schedule with deadlock

T_1	T_2
lock-X on A write (A)	
	lock-X on B write (B) wait for **lock-X** on A
wait for **lock-X** on B	

Deadlock Handling

- System is deadlocked if there is a set of transactions such that every transaction in the set is waiting for another transaction in the set.

- *Deadlock prevention* protocols ensure that the system will *never* enter into a deadlock state. Some prevention strategies :

  - Require that each transaction locks all its data items before it begins execution (predeclaration).

  - Impose partial ordering of all data items and require that a transaction can lock data items only in the order specified by the partial order (graph-based protocol).

More Deadlock Prevention Strategies

■ Following schemes use transaction timestamps for the sake of deadlock prevention alone.

■ **wait-die** scheme — non-preemptive

- older transaction may wait for younger one to release data item. Younger transactions never wait for older ones; they are rolled back instead.

- a transaction may die several times before acquiring needed data item

■ **wound-wait** scheme — preemptive

- older transaction *wounds* (forces rollback) of younger transaction instead of waiting for it. Younger transactions may wait for older ones.

- may be fewer rollbacks than *wait-die* scheme.

Deadlock prevention (Cont.)

■ Both in *wait-die* and in *wound-wait* schemes, a rolled back transactions is restarted with its original timestamp. Older transactions thus have precedence over newer ones, and starvation is hence avoided.

■ Timeout-Based Schemes:

- a transaction waits for a lock only for a specified amount of time. After that, the wait times out and the transaction is rolled back.

- thus deadlocks are not possible

- simple to implement; but starvation is possible. Also difficult to determine good value of the timeout interval.

Deadlock Detection

- Deadlocks can be described as a *wait-for graph*, which consists of a pair $G = (V,E)$,

    - V is a set of vertices (all the transactions in the system)

    - E is a set of edges; each element is an ordered pair $T_i \rightarrow T_j$.

- If $T_i \rightarrow T_j$ is in E, then there is a directed edge from T_i to T_j, implying that T_i is waiting for T_j to release a data item.

- When T_i requests a data item currently being held by T_j, then the edge $T_i \ T_j$ is inserted in the wait-for graph. This edge is removed only when T_j is no longer holding a data item needed by T_i.

- The system is in a deadlock state if and only if the wait-for graph has a cycle. Must invoke a deadlock-detection algorithm periodically to look for cycles.

Deadlock Detection (Cont.)

Wait-for graph without a cycle

Wait-for graph with a cycle

Deadlock Recovery

■ When deadlock is detected :

- Some transaction will have to rolled back (made a victim) to break deadlock. Select that transaction as victim that will incur minimum cost.

- Rollback -- determine how far to roll back transaction

  ▸ Total rollback: Abort the transaction and then restart it.

  ▸ More effective to roll back transaction only as far as necessary to break deadlock.

- Starvation happens if same transaction is always chosen as victim. Include the number of rollbacks in the cost factor to avoid starvation

Multiple Granularity

- Allow data items to be of various sizes and define a hierarchy of data granularities, where the small granularities are nested within larger ones

- Can be represented graphically as a tree (but don't confuse with tree-locking protocol)

- When a transaction locks a node in the tree *explicitly*, it *implicitly* locks all the node's descendents in the same mode.

- Granularity of locking (level in tree where locking is done):
  - **fine granularity** (lower in tree): high concurrency, high locking overhead
  - **coarse granularity** (higher in tree): low locking overhead, low concurrency

Example of Granularity Hierarchy

The levels, starting from the coarsest (top) level are

- *database*
- *area*
- *file*
- *record*

Intention Lock Modes

- In addition to S and X lock modes, there are three additional lock modes with multiple granularity:

    - **intention-shared** (IS): indicates explicit locking at a lower level of the tree but only with shared locks.

    - **intention-exclusive** (IX): indicates explicit locking at a lower level with exclusive or shared locks

    - **shared and intention-exclusive** (SIX): the subtree rooted by that node is locked explicitly in shared mode and explicit locking is being done at a lower level with exclusive-mode locks.

- intention locks allow a higher level node to be locked in S or X mode without having to check all descendent nodes.

Compatibility Matrix with Intention Lock Modes

■ The compatibility matrix for all lock modes is:

	IS	IX	S	SIX	X
IS	true	true	true	true	false
IX	true	true	false	false	false
S	true	false	true	false	false
SIX	true	false	false	false	false
X	false	false	false	false	false

Multiple Granularity Locking Scheme

- Transaction T_i can lock a node Q, using the following rules:
    1. The lock compatibility matrix must be observed.
    2. The root of the tree must be locked first, and may be locked in any mode.
    3. A node Q can be locked by T_i in S or IS mode only if the parent of Q is currently locked by T_i in either IX or IS mode.
    4. A node Q can be locked by T_i in X, SIX, or IX mode only if the parent of Q is currently locked by T_i in either IX or SIX mode.
    5. T_i can lock a node only if it has not previously unlocked any node (that is, T_i is two-phase).
    6. T_i can unlock a node Q only if none of the children of Q are currently locked by T_i.

- Observe that locks are acquired in root-to-leaf order, whereas they are released in leaf-to-root order.

- **Lock granularity escalation**: in case there are too many locks at a particular level, switch to higher granularity S or X lock

Timestamp-Based Protocols

■ Each transaction is issued a timestamp when it enters the system. If an old transaction T_i has time-stamp $TS(T_i)$, a new transaction T_j is assigned time-stamp $TS(T_j)$ such that $TS(T_i) < TS(T_j)$.

■ The protocol manages concurrent execution such that the time-stamps determine the serializability order.

■ In order to assure such behavior, the protocol maintains for each data Q two timestamp values:

- **W-timestamp**(Q) is the largest time-stamp of any transaction that executed **write**(Q) successfully.

- **R-timestamp**(Q) is the largest time-stamp of any transaction that executed **read**(Q) successfully.

Timestamp-Based Protocols (Cont.)

- The timestamp ordering protocol ensures that any conflicting **read** and **write** operations are executed in timestamp order.

- Suppose a transaction T_i issues a **read**(Q)

  1. If $TS(T_i) \leq$ **W**-timestamp(Q), then T_i needs to read a value of Q that was already overwritten.

     - Hence, the **read** operation is rejected, and T_i is rolled back.

  2. If $TS(T_i) \geq$ **W**-timestamp(Q), then the **read** operation is executed, and R-timestamp(Q) is set to **max**(R-timestamp(Q), $TS(T_i)$).

Timestamp-Based Protocols (Cont.)

- Suppose that transaction T_i issues **write**(Q).

  1. If $TS(T_i) <$ R-timestamp(Q), then the value of Q that T_i is producing was needed previously, and the system assumed that that value would never be produced.

     - Hence, the **write** operation is rejected, and T_i is rolled back.

  2. If $TS(T_i) <$ W-timestamp(Q), then T_i is attempting to write an obsolete value of Q.

     - Hence, this **write** operation is rejected, and T_i is rolled back.

  3. Otherwise, the **write** operation is executed, and W-timestamp(Q) is set to $TS(T_i)$.

Example Use of the Protocol

A partial schedule for several data items for transactions with timestamps 1, 2, 3, 4, 5

T_1	T_2	T_3	T_4	T_5
				read (X)
	read (Y)			
read (Y)				
		write (Y)		
		write (Z)		
				read (Z)
	read (Z)			
	abort			
read (X)				
			read (W)	
		write (W)		
		abort		
				write (Y)
				write (Z)

Correctness of Timestamp-Ordering Protocol

- The timestamp-ordering protocol guarantees serializability since all the arcs in the precedence graph are of the form:

  Thus, there will be no cycles in the precedence graph

- Timestamp protocol ensures freedom from deadlock as no transaction ever waits.

- But the schedule may not be cascade-free, and may not even be recoverable.

Recoverability and Cascade Freedom

- Problem with timestamp-ordering protocol:
  - Suppose T_i aborts, but T_j has read a data item written by T_i
  - Then T_j must abort; if T_j had been allowed to commit earlier, the schedule is not recoverable.
  - Further, any transaction that has read a data item written by T_j must abort
  - This can lead to cascading rollback --- that is, a chain of rollbacks
- Solution 1:
  - A transaction is structured such that its writes are all performed at the end of its processing
  - All writes of a transaction form an atomic action; no transaction may execute while a transaction is being written
  - A transaction that aborts is restarted with a new timestamp
- Solution 2: Limited form of locking: wait for data to be committed before reading it
- Solution 3: Use commit dependencies to ensure recoverability

Thomas' Write Rule

- Modified version of the timestamp-ordering protocol in which obsolete **write** operations may be ignored under certain circumstances.

- When T_i attempts to write data item Q, if $\text{TS}(T_i) < \text{W-timestamp}(Q)$, then T_i is attempting to write an obsolete value of $\{Q\}$.

  - Rather than rolling back T_i as the timestamp ordering protocol would have done, this $\{$**write**$\}$ operation can be ignored.

- Otherwise this protocol is the same as the timestamp ordering protocol.

- Thomas' Write Rule allows greater potential concurrency.

  - Allows some view-serializable schedules that are not conflict-serializable.

Validation-Based Protocol

- Execution of transaction T_i is done in three phases.

  1. **Read and execution phase**: Transaction T_i writes only to temporary local variables

  2. **Validation phase**: Transaction T_i performs a ``validation test'' to determine if local variables can be written without violating serializability.

  3. **Write phase**: If T_i is validated, the updates are applied to the database; otherwise, T_i is rolled back.

- The three phases of concurrently executing transactions can be interleaved, but each transaction must go through the three phases in that order.

  - Assume for simplicity that the validation and write phase occur together, atomically and serially
    - I.e., only one transaction executes validation/write at a time.

- Also called as **optimistic concurrency control** since transaction executes fully in the hope that all will go well during validation

Validation-Based Protocol (Cont.)

- Each transaction T_i has 3 timestamps
  - Start(T_i) : the time when T_i started its execution
  - Validation(T_i): the time when T_i entered its validation phase
  - Finish(T_i) : the time when T_i finished its write phase
- Serializability order is determined by timestamp given at validation time, to increase concurrency.
  - Thus TS(T_i) is given the value of Validation(T_i).
- This protocol is useful and gives greater degree of concurrency if probability of conflicts is low.
  - because the serializability order is not pre-decided, and
  - relatively few transactions will have to be rolled back.

Validation Test for Transaction T_j

- If for all T_i with TS (T_i) < TS (T_j) either one of the following condition holds:

  - **finish**(T_i) < **start**(T_j)

  - **start**(T_j) < **finish**(T_i) < **validation**(T_j) **and** the set of data items written by T_i does not intersect with the set of data items read by T_j.

  then validation succeeds and T_j can be committed. Otherwise, validation fails and T_j is aborted.

- *Justification*: Either the first condition is satisfied, and there is no overlapped execution, or the second condition is satisfied and

  - the writes of T_j do not affect reads of T_i since they occur after T_i has finished its reads.

  - the writes of T_i do not affect reads of T_j since T_j does not read any item written by T_i.

■ Example of schedule produced using validation

T_{25}	T_{26}
read (B)	
	read (B)
	$B := B$ 50
	read (A)
	$A := A + 50$
read (A)	
$\langle validate \rangle$	
display $(A + B)$	
	$\langle validate \rangle$
	write (B)
	write (A)

Multiversion Schemes

- Multiversion schemes keep old versions of data item to increase concurrency.

  - Multiversion Timestamp Ordering
  - Multiversion Two-Phase Locking

- Each successful **write** results in the creation of a new version of the data item written.

- Use timestamps to label versions.

- When a **read**(Q) operation is issued, select an appropriate version of Q based on the timestamp of the transaction, and return the value of the selected version.

- **read**s never have to wait as an appropriate version is returned immediately.

Multiversion Timestamp Ordering

- Each data item Q has a sequence of versions $<Q_1, Q_2,...., Q_m>$. Each version Q_k contains three data fields:

  - **Content** -- the value of version Q_k.

  - **W-timestamp**(Q_k) -- timestamp of the transaction that created (wrote) version Q_k

  - **R-timestamp**(Q_k) -- largest timestamp of a transaction that successfully read version Q_k

- when a transaction T_i creates a new version Q_k of Q, Q_k's W-timestamp and R-timestamp are initialized to TS(T_i).

- R-timestamp of Q_k is updated whenever a transaction T_j reads Q_k, and TS(T_j) > R-timestamp(Q_k).

Multiversion Timestamp Ordering (Cont)

- Suppose that transaction T_i issues a **read**(Q) or **write**(Q) operation. Let Q_k denote the version of Q whose write timestamp is the largest write timestamp less than or equal to TS(T_i).

    1. If transaction T_i issues a **read**(Q), then the value returned is the content of version Q_k.

    2. If transaction T_i issues a **write**(Q)

        1. if TS(T_i) < R-timestamp(Q_k), then transaction T_i is rolled back.

        2. if TS(T_i) = W-timestamp(Q_k), the contents of Q_k are overwritten

        3. else a new version of Q is created.

- Observe that

    - Reads always succeed

    - A write by T_i is rejected if some other transaction T_j that (in the serialization order defined by the timestamp values) should read T_i's write, has already read a version created by a transaction older than T_i.

- Protocol guarantees serializability

Multiversion Two-Phase Locking

■ Differentiates between read-only transactions and update transactions

■ *Update transactions* acquire read and write locks, and hold all locks up to the end of the transaction. That is, update transactions follow rigorous two-phase locking.

   ● Each successful **write** results in the creation of a new version of the data item written.

   ● each version of a data item has a single timestamp whose value is obtained from a counter **ts-counter** that is incremented during commit processing.

■ *Read-only transactions* are assigned a timestamp by reading the current value of **ts-counter** before they start execution; they follow the multiversion timestamp-ordering protocol for performing reads.

Multiversion Two-Phase Locking (Cont.)

- When an update transaction wants to read a data item:
  - it obtains a shared lock on it, and reads the latest version.
- When it wants to write an item
  - it obtains X lock on; it then creates a new version of the item and sets this version's timestamp to ∞.
- When update transaction T_i completes, commit processing occurs:
  - T_i sets timestamp on the versions it has created to **ts-counter** + 1
  - T_i increments **ts-counter** by 1
- Read-only transactions that start after T_i increments **ts-counter** will see the values updated by T_i.
- Read-only transactions that start before T_i increments the **ts-counter** will see the value before the updates by T_i.
- Only serializable schedules are produced.

MVCC: Implementation Issues

- Creation of multiple versions increases storage overhead

  - Extra tuples

  - Extra space in each tuple for storing version information

- Versions can, however, be garbage collected

  - E.g. if Q has two versions Q5 and Q9, and the oldest active transaction has timestamp > 9, than Q5 will never be required again

Snapshot Isolation

- Motivation: Decision support queries that read large amounts of data have concurrency conflicts with OLTP transactions that update a few rows
  - Poor performance results
- Solution 1: Give logical "snapshot" of database state to read only transactions, read-write transactions use normal locking
  - Multiversion 2-phase locking
  - Works well, but how does system know a transaction is read only?
- Solution 2: Give snapshot of database state to every transaction, updates alone use 2-phase locking to guard against concurrent updates
  - Problem: variety of anomalies such as lost update can result
  - Partial solution: snapshot isolation level (next slide)
    - Proposed by Berenson et al, SIGMOD 1995
    - Variants implemented in many database systems
      - E.g. Oracle, PostgreSQL, SQL Server 2005

Snapshot Isolation

- A transaction T1 executing with Snapshot Isolation

  - takes snapshot of committed data at start

  - always reads/modifies data in its own snapshot

  - updates of concurrent transactions are not visible to T1

  - writes of T1 complete when it commits

  - **First-committer-wins rule**:

    - Commits only if no other concurrent transaction has already written data that T1 intends to write.

T1	T2	T3
W(Y := 1) Commit		
	Start R(X) → 0 R(Y)→ 1	
		W(X:=2) W(Z:=3) Commit
	R(Z) → 0 R(Y) → 1 W(X:=3) Commit-Req Abort	

Concurrent updates not visible
Own updates are visible
Not first-committer of X
Serialization error, T2 is rolled back

Snapshot Read

■ Concurrent updates invisible to snapshot read

$X_0 = 100, \; Y_0 = 0$

T_1 deposits 50 in Y	T_2 withdraws 50 from X
$r_1(X_0, 100)$ $r_1(Y_0, 0)$	
	$r_2(Y_0, 0)$ $r_2(X_0, 100)$ $w_2(X_2, 50)$
$w_1(Y_1, 50)$ $r_1(X_0, 100)$ (update by T_2 not seen) $r_1(Y_1, 50)$ (can see its own updates)	
	$r_2(Y_0, 0)$ (update by T_1 not seen)

$X_2 = 50, \; Y_1 = 50$

Snapshot Write: First Committer Wins

- Variant: "**First-updater-wins**"
  - ▸ Check for concurrent updates when write occurs by locking item
    - – But lock should be held till all concurrent transactions have finished
  - ▸ (Oracle uses this plus some extra features)
  - ▸ Differs only in when abort occurs, otherwise equivalent

Benefits of SI

- Reading is *never* blocked,
  - and also doesn't block other txns activities
- Performance similar to Read Committed
- Avoids the usual anomalies
  - No dirty read
  - No lost update
  - No non-repeatable read
  - Predicate based selects are repeatable (no phantoms)
- Problems with SI
  - SI does not always give serializable executions
    - Serializable: among two concurrent txns, one sees the effects of the other
    - In SI: neither sees the effects of the other
  - Result: Integrity constraints can be violated

Snapshot Isolation

- E.g. of problem with SI
    - T1: x:=y
    - T2: y:= x
    - Initially x = 3 and y = 17
        - Serial execution: x = ??, y = ??
        - if both transactions start at the same time, with snapshot isolation: x = ?? , y = ??
- Called **skew write**
- Skew also occurs with inserts
    - E.g:
        - Find max order number among all orders
        - Create a new order with order number = previous max + 1

Snapshot Isolation Anomalies

- SI breaks serializability when txns modify *different* items, each based on a previous state of the item the other modified
  - Not very common in practice
    - E.g., the TPC-C benchmark runs correctly under SI
    - when txns conflict due to modifying different data, there is usually also a shared item they both modify too (like a total quantity) so SI will abort one of them
  - But does occur
    - Application developers should be careful about write skew
- SI can also cause a read-only transaction anomaly, where read-only transaction may see an inconsistent state even if updaters are serializable
  - We omit details
- Using snapshots to verify primary/foreign key integrity can lead to inconsistency
  - Integrity constraint checking usually done outside of snapshot

SI In Oracle and PostgreSQL

- **Warning**: SI used when isolation level is set to serializable, by Oracle, and PostgreSQL versions prior to 9.1

    - PostgreSQL's implementation of SI (versions prior to 9.1) described in Section 26.4.1.3

    - Oracle implements "first updater wins" rule (variant of "first committer wins")

        ‣ concurrent writer check is done at time of write, not at commit time

        ‣ Allows transactions to be rolled back earlier

        ‣ Oracle and PostgreSQL < 9.1 do not support true serializable execution

    - PostgreSQL 9.1 introduced new protocol called "Serializable Snapshot Isolation" (SSI)

        ‣ Which guarantees true serializabilty including handling predicate reads (coming up)

SI In Oracle and PostgreSQL

- Can sidestep SI for specific queries by using **select .. for update** in Oracle and PostgreSQL

    - E.g.,

        1. **select max**(orderno) **from** orders <u>**for update**</u>

        2. read value into local variable maxorder

        3. insert into orders (maxorder+1, …)

    - Select for update (SFU) treats all data read by the query as if it were also updated, preventing concurrent updates

    - Does not always ensure serializability since phantom phenomena can occur (coming up)

- In PostgreSQL versions < 9.1, SFU locks the data item, but releases locks when the transaction completes, even if other concurrent transactions are active

    - Not quite same as SFU in Oracle, which keeps locks until all

    - concurrent transactions have completed

Insert and Delete Operations

- If two-phase locking is used :
  - A **delete** operation may be performed only if the transaction deleting the tuple has an exclusive lock on the tuple to be deleted.
  - A transaction that inserts a new tuple into the database is given an X-mode lock on the tuple

- Insertions and deletions can lead to the **phantom phenomenon**.
  - A transaction that scans a relation
    - (e.g., find sum of balances of all accounts in Perryridge)

    and a transaction that inserts a tuple in the relation
    - (e.g., insert a new account at Perryridge)

    (conceptually) conflict in spite of not accessing any tuple in common.
  - If only tuple locks are used, non-serializable schedules can result
    - E.g. the scan transaction does not see the new account, but reads some other tuple written by the update transaction

Insert and Delete Operations (Cont.)

- The transaction scanning the relation is reading information that indicates what tuples the relation contains, while a transaction inserting a tuple updates the same information.
  - The conflict should be detected, e.g. by locking the information.
- One solution:
  - Associate a data item with the relation, to represent the information about what tuples the relation contains.
  - Transactions scanning the relation acquire a shared lock in the data item,
  - Transactions inserting or deleting a tuple acquire an exclusive lock on the data item. (Note: locks on the data item do not conflict with locks on individual tuples.)
- Above protocol provides very low concurrency for insertions/deletions.
- Index locking protocols provide higher concurrency while preventing the phantom phenomenon, by requiring locks on certain index buckets.

Index Locking Protocol

- Index locking protocol:
  - Every relation must have at least one index.
  - A transaction can access tuples only after finding them through one or more indices on the relation
  - A transaction T_i that performs a lookup must lock all the index leaf nodes that it accesses, in S-mode
    - Even if the leaf node does not contain any tuple satisfying the index lookup (e.g. for a range query, no tuple in a leaf is in the range)
  - A transaction T_i that inserts, updates or deletes a tuple t_i in a relation r
    - must update all indices to r
    - must obtain exclusive locks on all index leaf nodes affected by the insert/update/delete
  - The rules of the two-phase locking protocol must be observed
- Guarantees that phantom phenomenon won't occur

Next-Key Locking

- Index-locking protocol to prevent phantoms required locking entire leaf
    - Can result in poor concurrency if there are many inserts
- Alternative: for an index lookup
    - Lock all values that satisfy index lookup (match lookup value, or fall in lookup range)
    - Also lock next key value in index
    - Lock mode: S for lookups, X for insert/delete/update
- Ensures that range queries will conflict with inserts/deletes/updates
    - Regardless of which happens first, as long as both are concurrent

Concurrency in Index Structures

■ Indices are unlike other database items in that their only job is to help in accessing data.

■ Index-structures are typically accessed very often, much more than other database items.

  ● Treating index-structures like other database items, e.g. by 2-phase locking of index nodes can lead to low concurrency.

■ There are several index concurrency protocols where locks on internal nodes are released early, and not in a two-phase fashion.

  ● It is acceptable to have nonserializable concurrent access to an index as long as the accuracy of the index is maintained.

    ▸ In particular, the exact values read in an internal node of a B+-tree are irrelevant so long as we land up in the correct leaf node.

Concurrency in Index Structures (Cont.)

- Example of index concurrency protocol:

- Use **crabbing** instead of two-phase locking on the nodes of the B$^+$-tree, as follows. During search/insertion/deletion:
  - First lock the root node in shared mode.
  - After locking all required children of a node in shared mode, release the lock on the node.
  - During insertion/deletion, upgrade leaf node locks to exclusive mode.
  - When splitting or coalescing requires changes to a parent, lock the parent in exclusive mode.

- Above protocol can cause excessive deadlocks
  - Searches coming down the tree deadlock with updates going up the tree
  - Can abort and restart search, without affecting transaction

- Better protocols are available; see Section 16.9 for one such protocol, the B-link tree protocol
  - Intuition: release lock on parent before acquiring lock on child
    - And deal with changes that may have happened between lock release and acquire

Weak Levels of Consistency

- **Degree-two consistency:** differs from two-phase locking in that S-locks may be released at any time, and locks may be acquired at any time
    - X-locks must be held till end of transaction
    - Serializability is not guaranteed, programmer must ensure that no erroneous database state will occur]

- **Cursor stability**:
    - For reads, each tuple is locked, read, and lock is immediately released
    - X-locks are held till end of transaction
    - Special case of degree-two consistency

Weak Levels of Consistency in SQL

- **SQL allows non-serializable executions**
  - **Serializable:** is the default
  - **Repeatable read**: allows only committed records to be read, and repeating a read should return the same value (so read locks should be retained)
    - However, the phantom phenomenon need not be prevented
      - T1 may see some records inserted by T2, but may not see others inserted by T2
  - **Read committed**: same as degree two consistency, but most systems implement it as cursor-stability
  - **Read uncommitted**: allows even uncommitted data to be read
- **In many database systems, read committed is the default consistency level**
  - has to be explicitly changed to serializable when required
    - **set isolation level serializable**

Transactions across User Interaction

- Many applications need transaction support across user interactions
  - Can't use locking
  - Don't want to reserve database connection per user
- Application level concurrency control
  - Each tuple has a version number
  - Transaction notes version number when reading tuple
    - **select** r.balance, r.version **into** :A, :version
      from r **where** acctId =23
  - When writing tuple, check that current version number is same as the version when tuple was read
    - **update** r **set** r.balance = r.balance + :deposit
      where acctId = 23 **and** r.version = :version
- Equivalent to **optimistic concurrency control without validating read set**
- Used internally in Hibernate ORM system, and manually in many applications
- Version numbering can also be used to support first committer wins check of snapshot isolation
  - Unlike SI, reads are not guaranteed to be from a single snapshot

End of Chapter

Thanks to Alan Fekete and Sudhir Jorwekar for Snapshot Isolation examples

modified from:

Database System Concepts, 6th Ed.

Figure 15.01

	S	X
S	true	false
X	false	false

Figure 15.04

T_1	T_2	concurrency-control manager
lock-x (B)		
		grant-x (B, T_1)
read (B)		
$B := B - 50$		
write (B)		
unlock (B)		
	lock-s (A)	
		grant-s (A, T_2)
	read (A)	
	unlock (A)	
	lock-s (B)	
		grant-s (B, T_2)
	read (B)	
	unlock (B)	
	display $(A + B)$	
lock-x (A)		
		grant-x (A, T_2)
read (A)		
$A := A + 50$		
write (A)		
unlock (A)		

Figure 15.07

T_3	T_4
lock-x (B)	
read (B)	
B := B - 50	
write (B)	
	lock-s (A)
	read (A)
	lock-s (B)
lock-x (A)	

Figure 15.08

T_5	T_6	T_7
lock-x (A) read (A) lock-s (B) read (B) write (A) unlock (A)		
	lock-x (A) read (A) write (A) unlock (A)	
		lock-s (A) read (A)

Figure 15.09

T_8	T_9
lock-s (a_1)	
	lock-s (a_1)
lock-s (a_2)	
	lock-s (a_2)
lock-s (a_3)	
lock-s (a_4)	
	unlock-s (a_3)
	unlock-s (a_4)
lock-s (a_n)	
upgrade (a_1)	

Figure 15.10

Figure 15.11

Figure 15.12

T_{10}	T_{11}	T_{12}	T_{13}
lock-x (B)			
	lock-x (D)		
	lock-x (H)		
	unlock (D)		
lock-x (E)			
lock-x (D)			
unlock (B)			
unlock (E)			
		lock-x (B)	
		lock-x (E)	
	unlock (H)		
lock-x (G)			
unlock (D)			
			lock-x (D)
			lock-x (H)
			unlock (D)
			unlock (H)
		unlock (E)	
		unlock (B)	
unlock (G)			

Figure 15.13

Figure 15.14

Figure 15.15

Figure 15.16

	IS	IX	S	SIX	X
IS	true	true	true	true	false
IX	true	true	false	false	false
S	true	false	true	false	false
SIX	true	false	false	false	false
X	false	false	false	false	false

Figure 15.17

T_{25}	T_{26}
read (B)	
	read (B)
	$B := B - 50$
	write (B)
read (A)	
	read (A)
display $(A + B)$	
	$A := A + 50$
	write (A)
	display $(A + B)$

Figure 15.18

T_{27}	T_{28}
read (Q)	
	write (Q)
write (Q)	

Figure 15.19

T_{25}	T_{26}
read (B)	
	read (B)
	$B := B \; 50$
	read (A)
	$A := A + 50$
read (A)	
$\langle \text{validate} \rangle$	
display $(A + B)$	
	$\langle \text{validate} \rangle$
	write (B)
	write (A)

Figure 15.20

T_{32}	T_{33}
lock-s (Q)	
read (Q)	
unlock (Q)	
	lock-x (Q)
	read (Q)
	write (Q)
	unlock (Q)
lock-s (Q)	
read (Q)	
unlock (Q)	

Figure 15.21

Figure 15.22
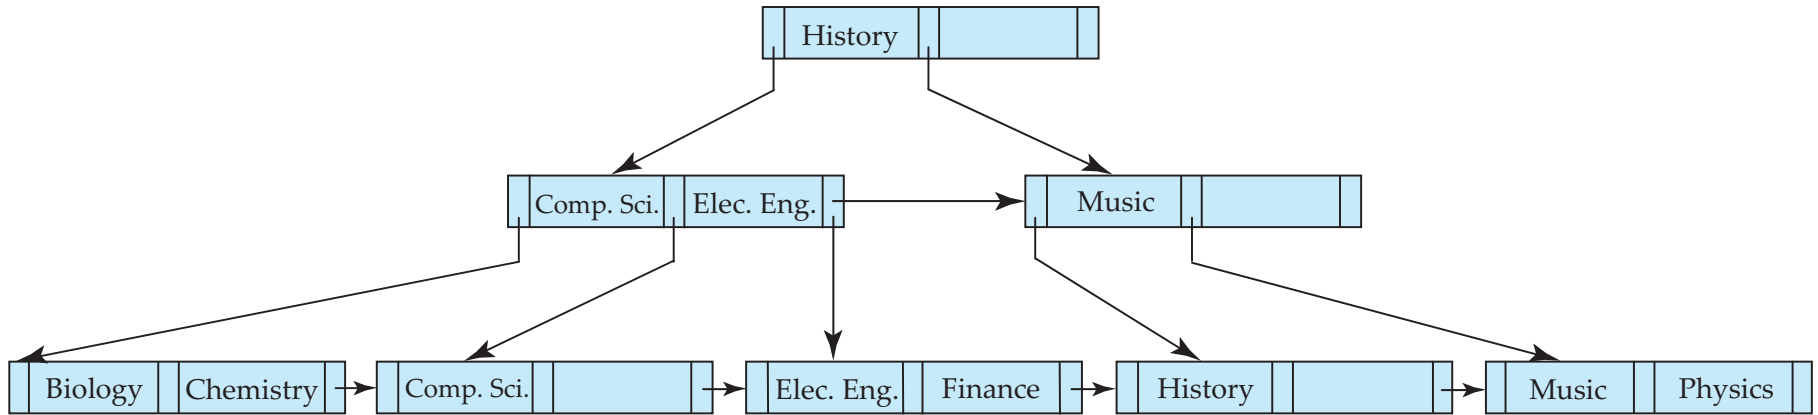
Figure 15.23

	S	X	I
S	true	false	false
X	false	false	false
I	false	false	true

Figure in-15.1

T_{27}	T_{28}	T_{29}
read (Q)		
	write (Q)	
write (Q)		
		write (Q)

Chapter 11: Indexing and Storage

Chapter 11: Indexing and Storage

- **DBMS Storage**
  - Memory hierarchy
  - File Organization
  - Buffering

- **Indexing**
  - Basic Concepts
  - B$^+$-Trees
  - Static Hashing
  - Index Definition in SQL
  - Multiple-Key Access

Memory Hierarchy

DBMS Storage

- Modern Computers have different types of memory

  - Cache, Main Memory, Harddisk, SSD, …

- Memory types have different characteristics in terms of

  - Persistent vs. volatile

  - Speed (random vs. sequential access)

  - Size

  - Price – this usually determines size

- Database systems are designed to be use these different memory types effectively

  - Need for persistent storage: the state of the database needs to be written to persistent storage

    - guarantee database content is not lost when the computer is shutdown

  - Moving data between different types of memory

    - Want to use fast memory to speed-up operations

    - Need slower memory for the size

Storage Hierarchy

Persistent storage

Size

Speed

cache

main memory

flash memory

magnetic disk

optical disk

magnetic tapes
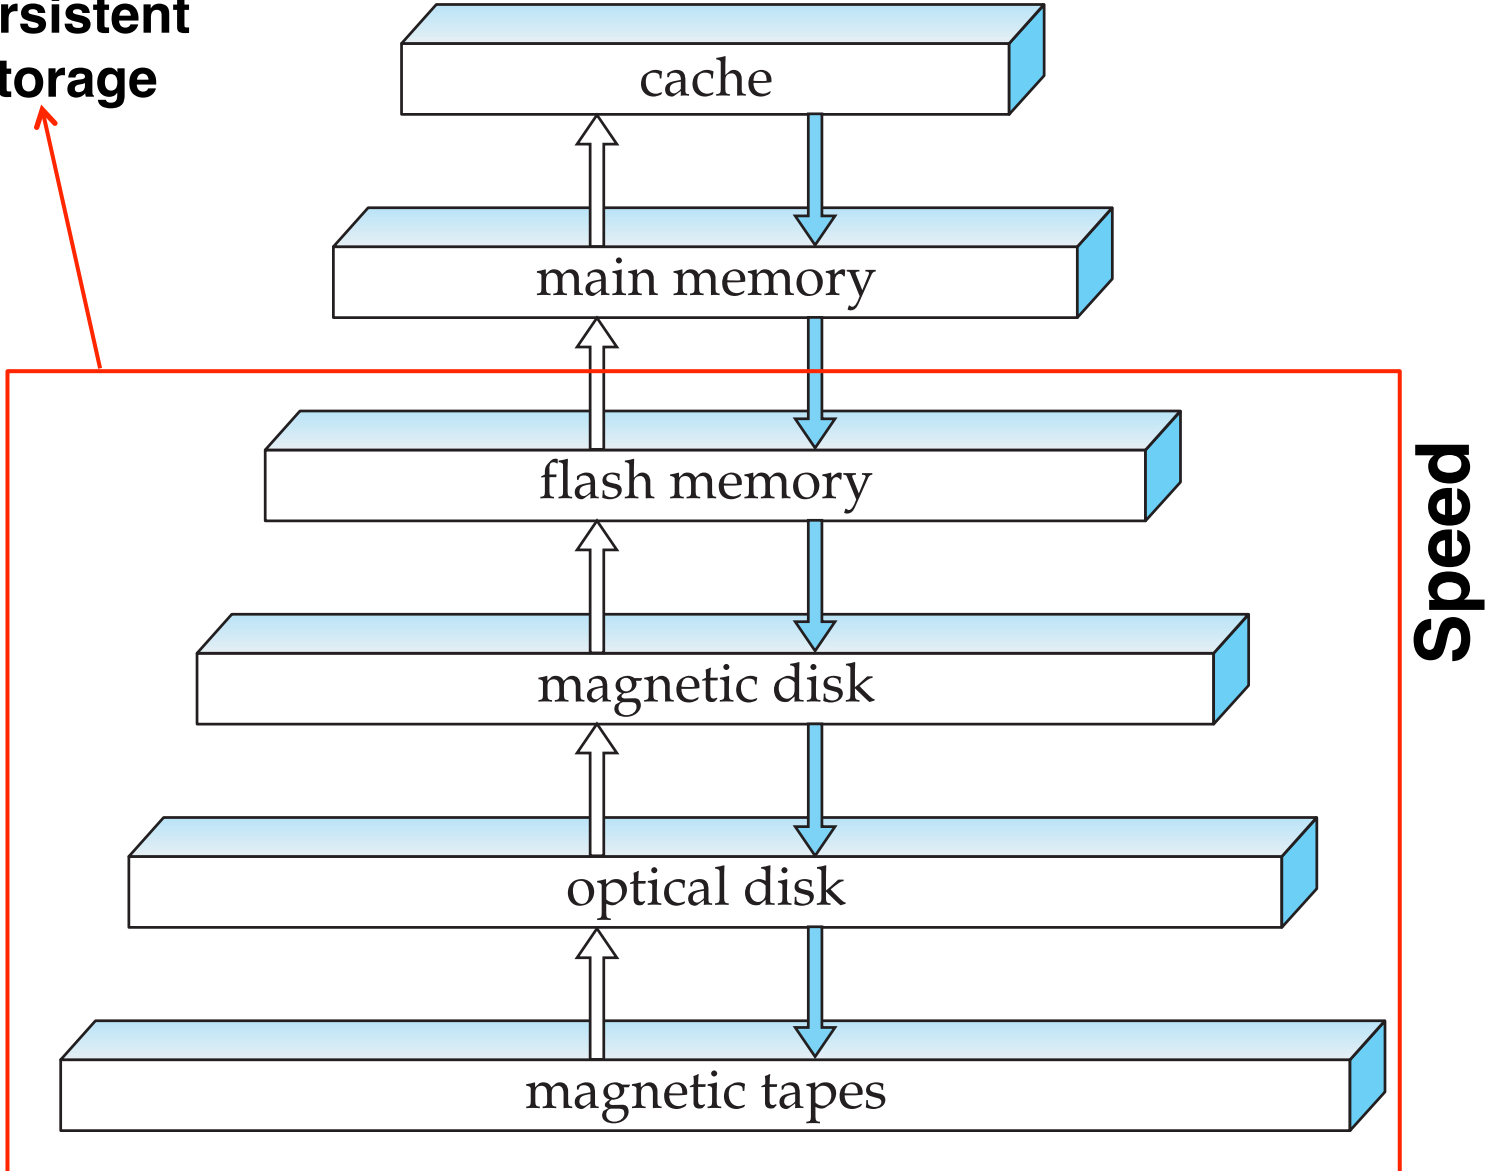
Main Memory vs. Disk

- Why do we not only use main memory
  - What if database does not fit into main memory?
  - Main memory is volatile
- Main memory vs. disk
  - Given available main memory when do we keep which part of the database in main memory
    - **Buffer manager**: Component of DBMS that decides when to move data between disk and main memory
  - How do we ensure transaction property durability
    - Buffer manager needs to make sure data written by committed transactions is written to disk to ensure durability

Magnetic Hard Disk Mechanism

NOTE: Diagram is schematic, and simplifies the structure of actual disk drives

Performance Measures of Disks

- **Access time** – the time it takes from when a read or write request is issued to when data transfer begins. Consists of:
  - **Seek time** – time it takes to reposition the arm over the correct track.
    - Average seek time is 1/2 the worst case seek time.
      - Would be 1/3 if all tracks had the same number of sectors, and we ignore the time to start and stop arm movement
    - 4 to 10 milliseconds on typical disks
  - **Rotational latency** – time it takes for the sector to be accessed to appear under the head.
    - Average latency is 1/2 of the worst case latency.
    - 4 to 11 milliseconds on typical disks (5400 to 15000 r.p.m.)
- **Data-transfer rate** – the rate at which data can be retrieved from or stored to the disk.
  - 25 to 100 MB per second max rate, lower for inner tracks
  - Multiple disks may share a controller, so rate that controller can handle is also important
    - E.g. SATA: 150 MB/sec, SATA-II 3Gb (300 MB/sec)
    - Ultra 320 SCSI: 320 MB/s, SAS (3 to 6 Gb/sec)
    - Fiber Channel (FC2Gb or 4Gb): 256 to 512 MB/s

Random vs. Sequential Access

- Transfer of data from disk has a minimal size = 1 block
    - Reading 1 byte is as fast as reading one block (e.g., 4KB)
- **Random Access**
    - Read data from anywhere on the disk
    - Need to get to the right track (**seek time**)
    - Need to wait until the right sector is under the arm (on avg ½ time for one rotation) (**rotational delay**)
    - Then can transfer data at ~ **transfer rate**
- **Sequential Access**
    - Read data that is on the current track + sector
    - can transfer data at ~ **transfer rate**
- Reading large number of small pieces of data randomly is very slow compared to sequential access
    - Thus, try layout data on disk in a way that enables sequential access

File Organization

File Organization

- The database is stored as a collection of **files**. Each file stores **records** (tuples from a table). A record is a sequence of **fields** (the attributes of a tuple).

- Reading one record of a time from disk would be very slow (random access)

  - Organize our database files in pages (size of block or larger)

  - Read/write data in units of pages

  - One page will usually contain several records

- One approach:

  - assume record size is fixed

  - each file has records of one particular type only

  - different files are used for different relations

  This case is easiest to implement; will consider variable length records later.

Fixed-Length Records

■ Simple approach:

- Store record i starting from byte $n * (i – 1)$, where n is the size of each record. Put maximal P / n records on each page.

- Record access is simple but records may cross blocks

  ▸ Modification: do not allow records to cross block boundaries

■ Deletion of record i: alternatives:

- move records $i + 1, . . ., n$ to $i, . . . , n – 1$

- move record n to i

- do not move records, but link all free records on a *free list*

record 0	10101	Srinivasan	Comp. Sci.	65000
record 1	12121	Wu	Finance	90000
record 2	15151	Mozart	Music	40000
record 3	22222	Einstein	Physics	95000
record 4	32343	El Said	History	60000
record 5	33456	Gold	Physics	87000
record 6	45565	Katz	Comp. Sci.	75000
record 7	58583	Califieri	History	62000
record 8	76543	Singh	Finance	80000
record 9	76766	Crick	Biology	72000
record 10	83821	Brandt	Comp. Sci.	92000
record 11	98345	Kim	Elec. Eng.	80000

Free Lists

- Store the address of the first deleted record in the file header.

- Use this first record to store the address of the second deleted record, and so on

- Can think of these stored addresses as pointers since they "point" to the location of a record.

header				
record 0	10101	Srinivasan	Comp. Sci.	65000
record 1				
record 2	15151	Mozart	Music	40000
record 3	22222	Einstein	Physics	95000
record 4				
record 5	33456	Gold	Physics	87000
record 6				
record 7	58583	Califieri	History	62000
record 8	76543	Singh	Finance	80000
record 9	76766	Crick	Biology	72000
record 10	83821	Brandt	Comp. Sci.	92000
record 11	98345	Kim	Elec. Eng.	80000

Variable-Length Records

- Variable-length records arise in database systems in several ways:
  - Storage of multiple record types in a file.
  - Record types that allow variable lengths for one or more fields such as strings (**varchar**)
  - Record types that allow repeating fields (used in some older data models).
- Attributes are stored in order
- Variable length attributes represented by fixed size (offset, length), with actual data stored after all fixed length attributes
- Null values represented by null-value bitmap

Null bitmap (stored in 1 byte)

0000

21, 5	26, 10	36, 10	65000		10101	Srinivasan	Comp. Sci.

Bytes 0   4   8   12   20 21   26   36   45

Variable-Length Records: Slotted Page Structure

- **Slotted page** header contains:
  - number of record entries
  - end of free space in the block
  - location and size of each record

- Records can be moved around within a page to keep them contiguous with no empty space between them; entry in the header must be updated.

- Pointers should not point directly to record — instead they should point to the entry for the record in header.

Organization of Records in Files

- **Heap** – a record can be placed anywhere in the file where there is space

    - Deletion efficient

    - Insertion efficient

    - Search is expensive

        - Example: Get instructor with name Glavic

            – Have to search through all instructors

- **Sequential** – store records in sequential order, based on the value of some search key of each record

    - Deletion expensive and/or waste of space

    - Insertion expensive and/or waste of space

    - Search is efficient (e.g., binary search)

        - As long as the search is on the search key we are ordering on

Buffering

Buffer Manager

- Buffer Manager
  - Responsible for loading pages from disk and writing modified pages back to disk

- Handling blocks
  1. If the block is already in the buffer, the buffer manager returns the address of the block in main memory
  2. If the block is not in the buffer, the buffer manager
     1. Allocates space in the buffer for the block
        1. Replacing (throwing out) some other block, if required, to make space for the new block.
        2. Replaced block written back to disk only if it was modified since the most recent time that it was written to/fetched from the disk.
     2. Reads the block from the disk to the buffer, and returns the address of the block in main memory to requester.

Buffer-Replacement Policies

- Most operating systems replace the block **least recently used** (LRU strategy)

- Idea behind LRU – use past pattern of block references as a predictor of future references

- Queries have well-defined access patterns (such as sequential scans), and a database system can use the information in a user's query to predict future references

  - LRU can be a bad strategy for certain access patterns involving repeated scans of data

    - For example: when computing the join of 2 relations r and s by a nested loops
      for each tuple *tr* of *r* do
        for each tuple *ts* of *s* do
          if the tuples *tr* and *ts* match …

  - Mixed strategy with hints on replacement strategy provided by the query optimizer is preferable

Buffer-Replacement Policies (Cont.)

- **Pinned block** – memory block that is not allowed to be written back to disk. E.g., an operation still needs this block.

- **Toss-immediate** strategy – frees the space occupied by a block as soon as the final tuple of that block has been processed

- **Most recently used (MRU) strategy** – system must pin the block currently being processed. After the final tuple of that block has been processed, the block is unpinned, and it becomes the most recently used block.

- Buffer manager can use statistical information regarding the probability that a request will reference a particular relation
  - E.g., the data dictionary is frequently accessed. Heuristic: keep data-dictionary blocks in main memory buffer

- Buffer managers also support **forced output** of blocks for the purpose of recovery (more in Chapter 16 in the textbook)

Indexing and Hashing

Basic Concepts

- Indexing mechanisms used to speed up access to desired data.

  - E.g., author catalog in library

- **Search Key** - attribute or set of attributes used to look up records in a file.

- An **index file** consists of records (called **index entries**) of the form

search-key	pointer

- Index files are typically much smaller than the original file

- Two basic kinds of indices:

  - **Ordered indices:** search keys are stored in some sorted order

  - **Hash indices:** search keys are distributed uniformly across "buckets" using a "hash function".

Index Evaluation Metrics

- Access types supported efficiently. E.g.,
  - records with a specified value in the attribute
  - or records with an attribute value falling in a specified range of values.
- Access time
- Insertion time
- Deletion time
- Space overhead

Ordered Indices

- In an **ordered index,** index entries are stored sorted on the search key value. E.g., author catalog in library.

- **Primary index:** in a sequentially ordered file, the index whose search key specifies the sequential order of the file.
  - Also called **clustering index**
  - The search key of a primary index is usually but not necessarily the primary key.

- **Secondary index**: an index whose search key specifies an order different from the sequential order of the file. Also called non-clustering index**.**

- Index-sequential file**:** ordered sequential file with a primary index.

Secondary Indices Example
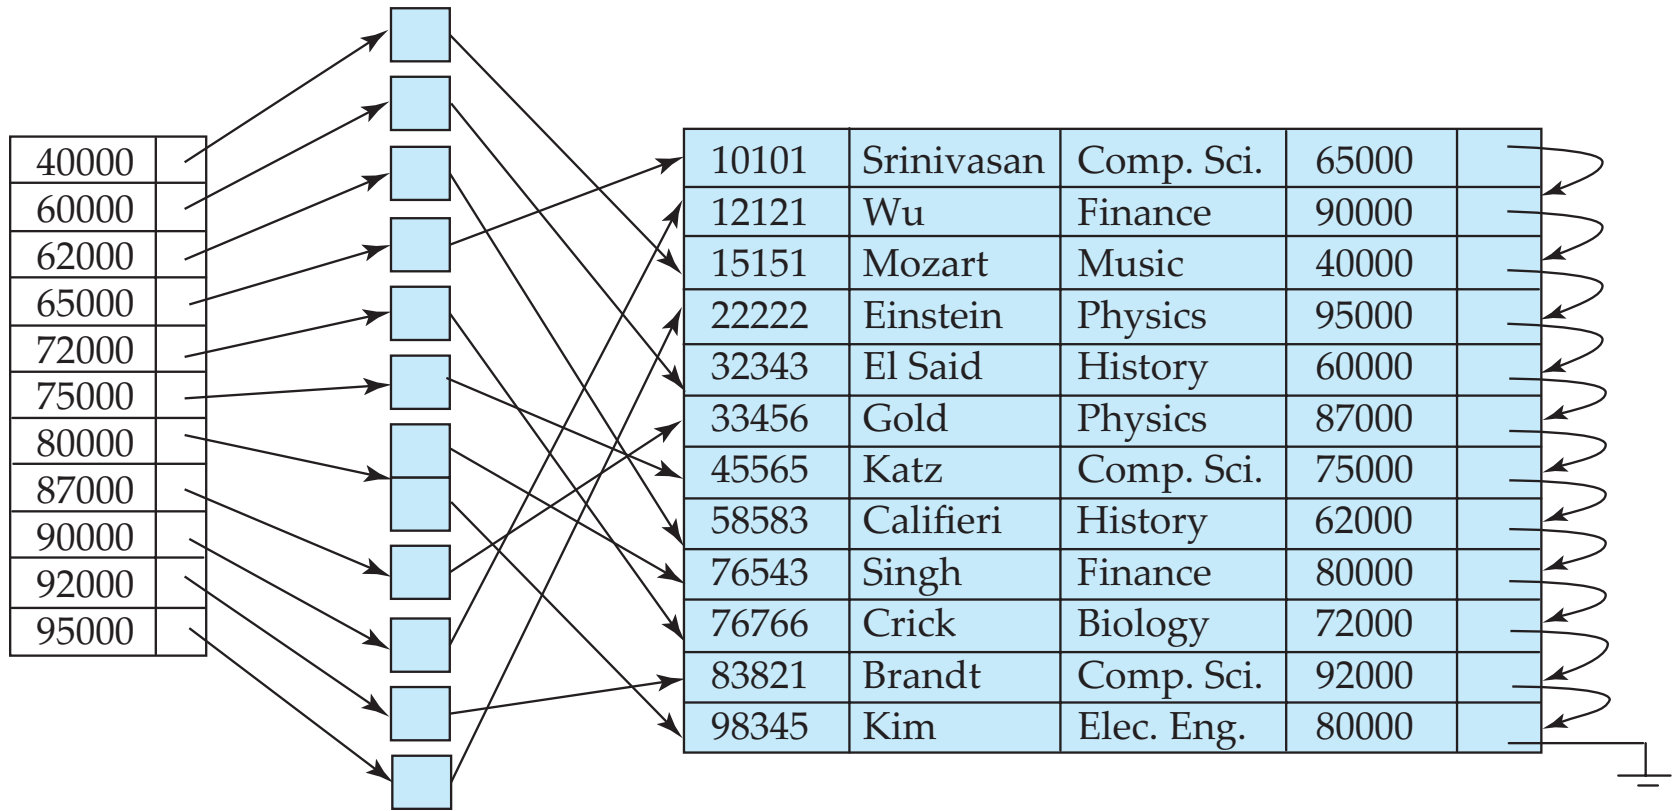
Secondary index on *salary* field of *instructor*

- Index record points to a bucket that contains pointers to all the actual records with that particular search-key value.

- Secondary indices have to be dense

Primary and Secondary Indices

■ Indices offer substantial benefits when searching for records.

■ BUT: Updating indices imposes overhead on database modification --when a file is modified, every index on the file must be updated,

■ Sequential scan using primary index is efficient, but a sequential scan using a secondary index is expensive

  ● Each record access may fetch a new block from disk

  ● Block fetch requires about 5 to 10 milliseconds, versus about 100 nanoseconds for memory access

Multilevel Index

- If primary index does not fit in memory, access becomes expensive.

- Solution: treat primary index kept on disk as a sequential file and construct a sparse index on it.
  - outer index – a sparse index of primary index
  - inner index – the primary index file

- If even outer index is too large to fit in main memory, yet another level of index can be created, and so on.

- Indices at all levels must be updated on insertion or deletion from the file.

outer index

index block 0

index block 1

inner index

data block 0

data block 1

Index Update: Deletion

10101	
32343	
76766	

10101	Srinivasan	Comp. Sci.	65000	
12121	Wu	Finance	90000	
15151	Mozart	Music	40000	
22222	Einstein	Physics	95000	
32343	El Said	History	60000	
33456	Gold	Physics	87000	
45565	Katz	Comp. Sci.	75000	
58583	Califieri	History	62000	
76543	Singh	Finance	80000	
76766	Crick	Biology	72000	
83821	Brandt	Comp. Sci.	92000	
98345	Kim	Elec. Eng.	80000	
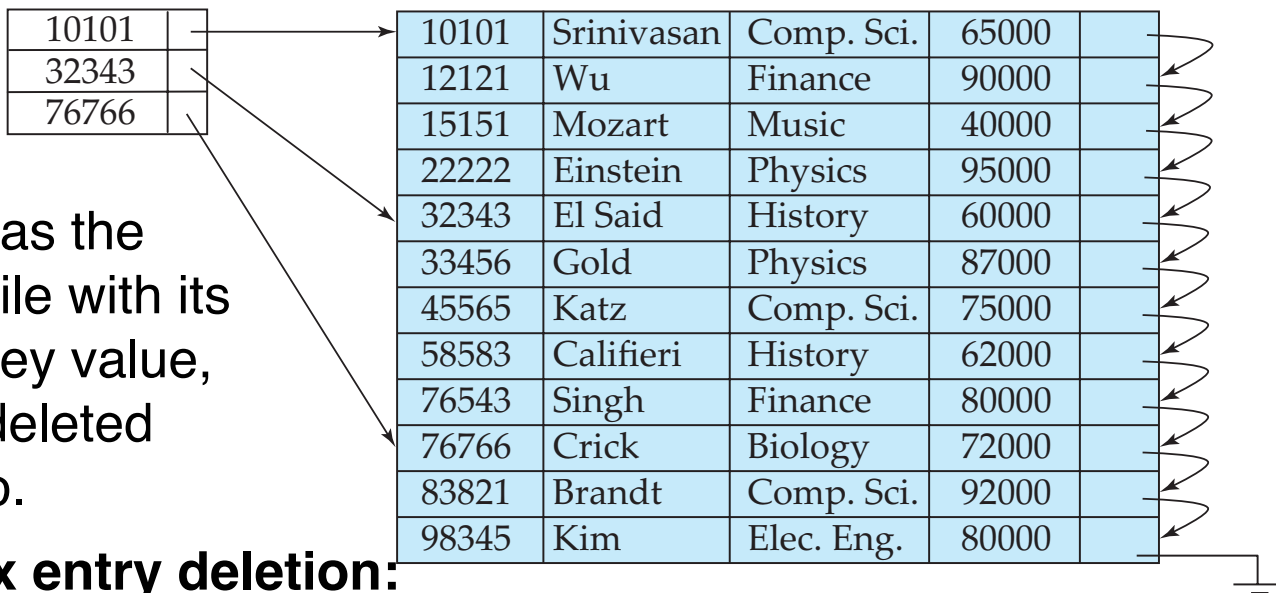
■ If deleted record was the only record in the file with its particular search-key value, the search-key is deleted from the index also.

■ **Single-level index entry deletion:**

- **Dense indices** – deletion of search-key is similar to file record deletion.

- **Sparse indices** –

    ▸ if an entry for the search key exists in the index, it is deleted by replacing the entry in the index with the next search-key value in the file (in search-key order).

    ▸ If the next search-key value already has an index entry, the entry is deleted instead of being replaced.

Index Update: Insertion

■ **Single-level index insertion:**

- Perform a lookup using the search-key value appearing in the record to be inserted.

- **Dense indices** – if the search-key value does not appear in the index, insert it.

- **Sparse indices** – if index stores an entry for each block of the file, no change needs to be made to the index unless a new block is created.

  ▸ If a new block is created, the first search-key value appearing in the new block is inserted into the index.

■ **Multilevel insertion and deletion:** algorithms are simple extensions of the single-level algorithms

Secondary Indices

- Frequently, one wants to find all the records whose values in a certain field (which is not the search-key of the primary index) satisfy some condition.

  - Example 1: In the *instructor* relation stored sequentially by ID, we may want to find all instructors in a particular department

  - Example 2: as above, but where we want to find all instructors with a specified salary or with salary in a specified range of values

- We can have a secondary index with an index record for each search-key value

B⁺-Tree Index

B⁺-tree indices are an alternative to indexed-sequential files.

- Disadvantage of indexed-sequential files
  - performance degrades as file grows, since many overflow blocks get created.
  - Periodic reorganization of entire file is required.
- Advantage of B⁺-tree index files:
  - automatically reorganizes itself with small, local, changes, in the face of insertions and deletions.
  - Reorganization of entire file is not required to maintain performance.
- (Minor) disadvantage of B⁺-trees:
  - extra insertion and deletion overhead, space overhead.
- Advantages of B⁺-trees outweigh disadvantages
  - B⁺-trees are used extensively

Example of B⁺-Tree
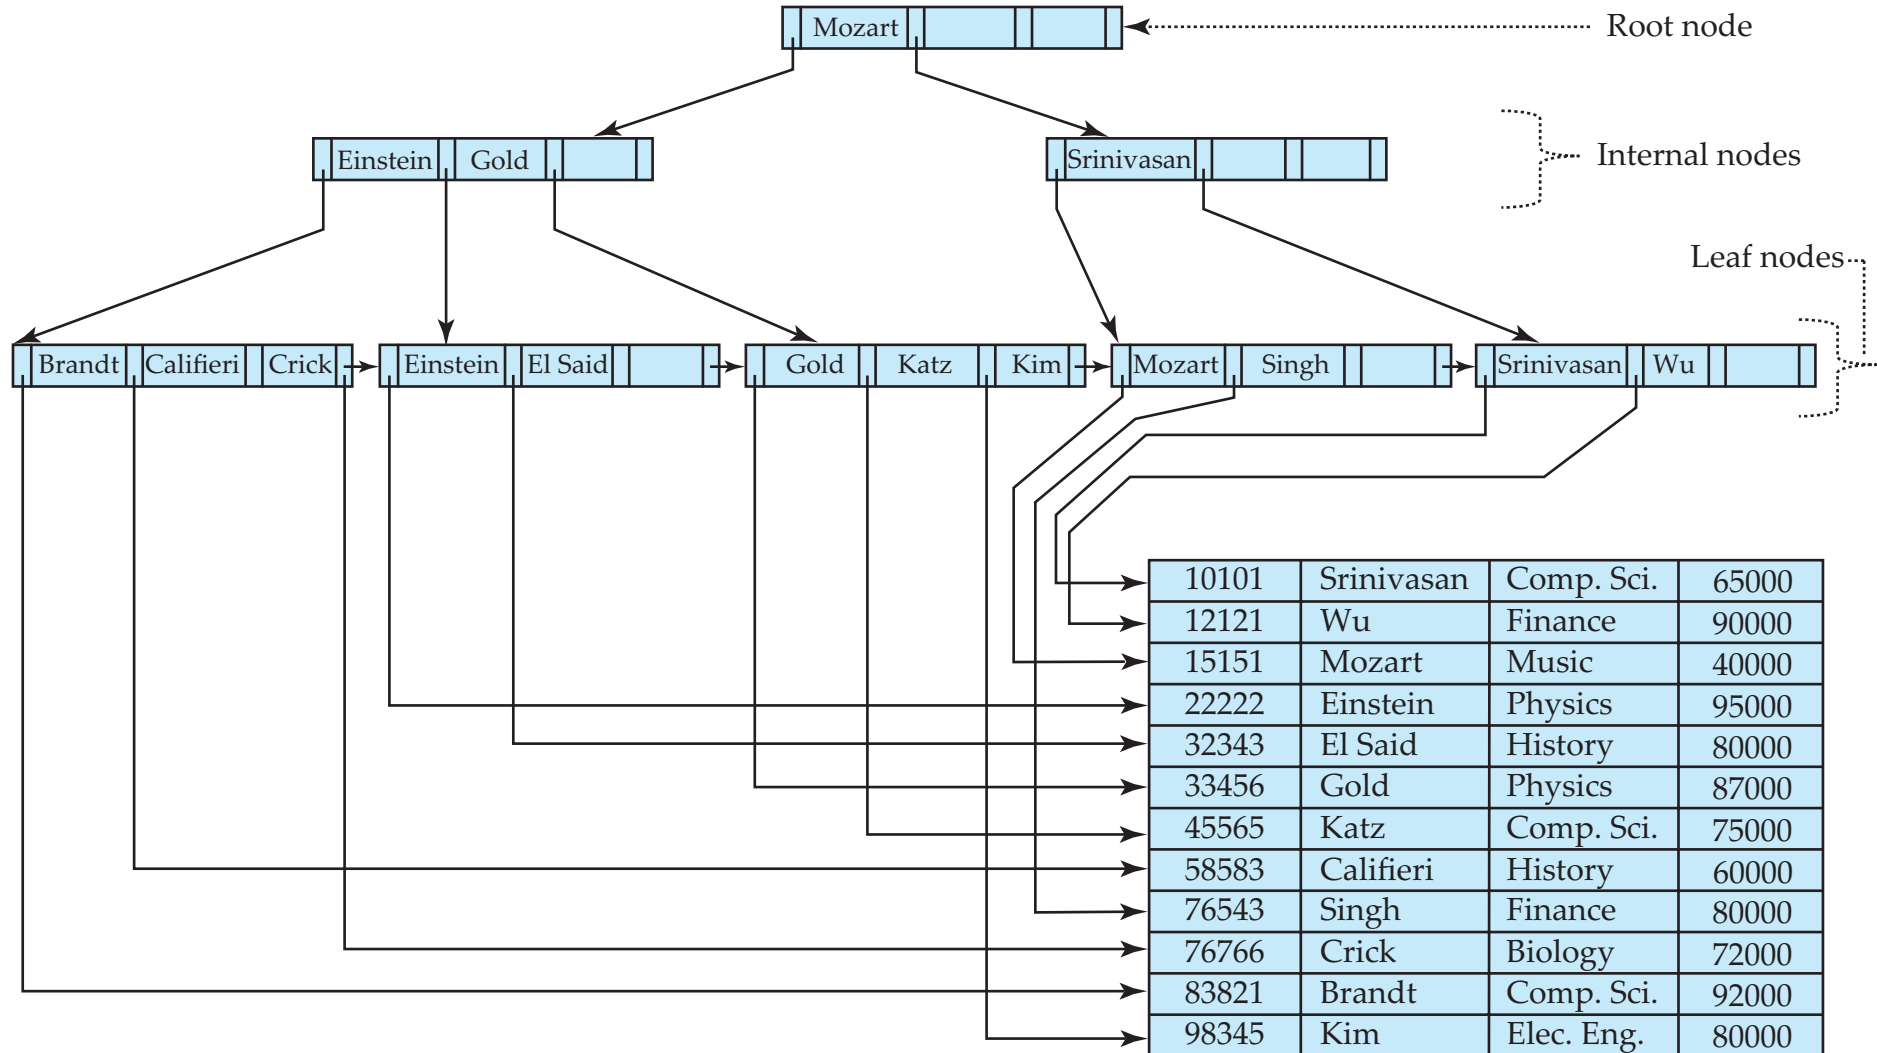
B+-Tree Index Files (Cont.)

A B+-tree is a rooted tree satisfying the following properties:

- All paths from root to leaf are of the same length

- Each node that is not a root or a leaf has between $\lceil n/2 \rceil$ and n children.

- A leaf node has between $\lceil (n-1)/2 \rceil$ and $n-1$ values

- Special cases:
  - If the root is not a leaf, it has at least 2 children.
  - If the root is a leaf (that is, there are no other nodes in the tree), it can have between 0 and ($n-1$) values.

B⁺-Tree Node Structure

- Typical node

P_1	K_1	P_2	$\ldots$	P_{n-1}	K_{n-1}	P_n

- K_i are the search-key values

- P_i are pointers to children (for non-leaf nodes) or pointers to records or buckets of records (for leaf nodes).

- The search-keys in a node are ordered

$$K_1 < K_2 < K_3 < \ldots < K_{n-1}$$

(Initially assume no duplicate keys, address duplicates later)

Leaf Nodes in B⁺-Trees

Properties of a leaf node:

- For $i = 1, 2, \ldots, n-1$, pointer P_i points to a file record with search-key value K_i,

- If L_i, L_j are leaf nodes and $i < j$, L_i's search-key values are less than or equal to L_j's search-key values

- P_n points to next leaf node in search-key order

leaf node

| | Brandt | | Califieri | | Crick | | → Pointer to next leaf node

10101	Srinivasan	Comp. Sci.	65000
12121	Wu	Finance	90000
15151	Mozart	Music	40000
22222	Einstein	Physics	95000
32343	El Said	History	80000
33456	Gold	Physics	87000
45565	Katz	Comp. Sci.	75000
58583	Califieri	History	60000
76543	Singh	Finance	80000
76766	Crick	Biology	72000
83821	Brandt	Comp. Sci.	92000
98345	Kim	Elec. Eng.	80000

Non-Leaf Nodes in B⁺-Trees

- Non leaf nodes form a multi-level sparse index on the leaf nodes. For a non-leaf node with m pointers:

  - All the search-keys in the subtree to which P_1 points are less than K_1

  - For $2 \leq i \leq n - 1$, all the search-keys in the subtree to which P_i points have values greater than or equal to K_{i-1} and less than K_i

  - All the search-keys in the subtree to which P_n points have values greater than or equal to K_{n-1}

P_1	K_1	P_2	...	P_{n-1}	K_{n-1}	P_n

Example of B⁺-tree

B⁺-tree for *instructor* file ($n = 6$)

- Leaf nodes must have between 3 and 5 values ($\lceil(n-1)/2\rceil$ and $n-1$, with $n = 6$).

- Non-leaf nodes other than root must have between 3 and 6 children ($\lceil(n/2)\rceil$ and n with $n = 6$).

- Root must have at least 2 children.

Observations about B⁺-trees

- Since the inter-node connections are done by pointers, "logically" close blocks need not be "physically" close.

- The non-leaf levels of the B⁺-tree form a hierarchy of sparse indices.

- The B⁺-tree contains a relatively small number of levels
    - Level below root has at least 2* $\lceil n/2 \rceil$ values
    - Next level has at least 2* $\lceil n/2 \rceil$ * $\lceil n/2 \rceil$ values
    - .. etc.

  - If there are K search-key values in the file, the tree height is no more than $\lceil \log_{\lceil n/2 \rceil}(K) \rceil$
  - thus searches can be conducted efficiently.

- Insertions and deletions to the main file can be handled efficiently, as the index can be restructured in logarithmic time (as we shall see).

Queries on B⁺-Trees

- Find record with search-key value *V.*
  1. *C=root*
  2. While C is not a leaf node {
     1. Let *i* be least value s.t. $V \leq K_i$.
     2. If no such exists, set *C = last non-null pointer in C*
     3. Else { if ($V = K_i$) Set $C = P_{i+1}$ else set $C = P_i$}
     }
  3. Let *i* be least value s.t. $K_i = V$
  4. If there is such a value *i*, follow pointer P_i to the desired record.
  5. Else no record with search-key value *k* exists.
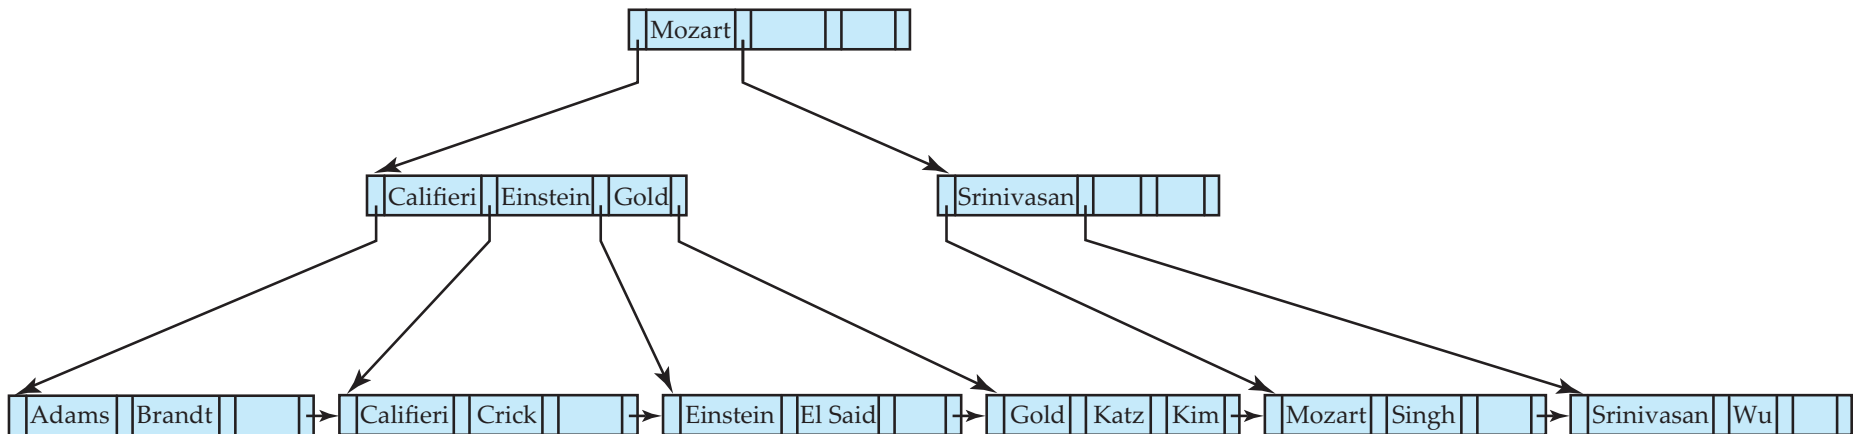
Handling Duplicates

- With duplicate search keys
  - In both leaf and internal nodes,
    - we cannot guarantee that $K_1 < K_2 < K_3 < \ldots < K_{n-1}$
    - but can guarantee $K_1 \leq K_2 \leq K_3 \leq \ldots \leq K_{n-1}$
  - Search-keys in the subtree to which P_i points
    - are $\leq K_i$, but not necessarily $< K_i$
    - To see why, suppose same search key value V is present in two leaf node L_i and L_{i+1}. Then in parent node K_i must be equal to V

Handling Duplicates

■ We modify find procedure as follows

- traverse P_i even if $V = K_i$

- As soon as we reach a leaf node C check if C has only search key values less than V

  ▸ if so set C = right sibling of C before checking whether C contains V

■ Procedure printAll

- uses modified find procedure to find first occurrence of V

- Traverse through consecutive leaves to find all occurrences of V

** Errata note: modified find procedure missing in first printing of 6th edition

Queries on B+-Trees (Cont.)

- If there are *K* search-key values in the file, the height of the tree is no more than $\lceil \log_{\lceil n/2 \rceil}(K) \rceil$.

- A node is generally the same size as a disk block, typically 4 kilobytes

  - and *n* is typically around 100 (40 bytes per index entry).

- With 1 million search key values and *n* = 100

  - at most $\log_{50}(1,000,000) = 4$ nodes are accessed in a lookup.

- Contrast this with a balanced binary tree with 1 million search key values — around 20 nodes are accessed in a lookup

  - above difference is significant since every node access may need a disk I/O, costing around 20 milliseconds

Updates on B+-Trees: Insertion

1. Find the leaf node in which the search-key value would appear

2. If the search-key value is already present in the leaf node

   1. Add record to the file

   2. If necessary add a pointer to the bucket.

3. If the search-key value is not present, then

   1. add the record to the main file (and create a bucket if necessary)

   2. If there is room in the leaf node, insert (key-value, pointer) pair in the leaf node

   3. Otherwise, split the node (along with the new (key-value, pointer) entry) as discussed in the next slide.

Updates on B+-Trees: Insertion (Cont.)

- Splitting a leaf node:
  - take the n (search-key value, pointer) pairs (including the one being inserted) in sorted order. Place the first $\lceil n/2 \rceil$ in the original node, and the rest in a new node.
  - let the new node be p, and let k be the least key value in p. Insert (k,p) in the parent of the node being split.
  - If the parent is full, split it and **propagate** the split further up.
- Splitting of nodes proceeds upwards till a node that is not full is found.
  - In the worst case the root node may be split increasing the height of the tree by 1.

Result of splitting node containing Brandt, Califieri and Crick on inserting Adams
Next step: insert entry with (Califieri,pointer-to-new-node) into parent

B⁺-Tree Insertion
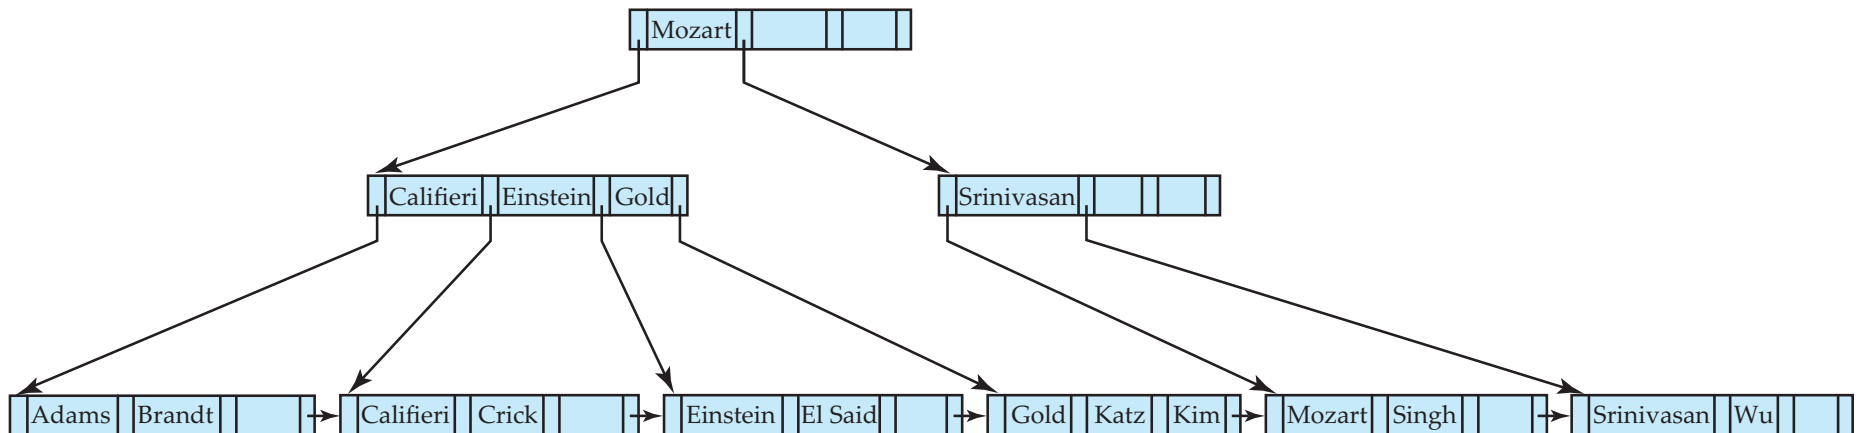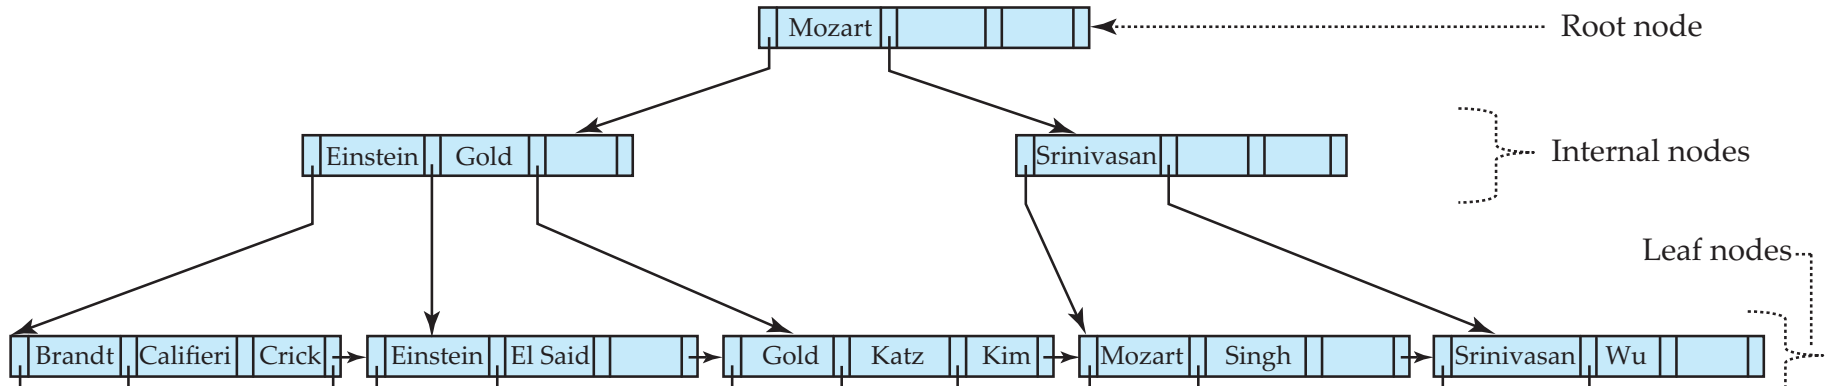
B⁺-Tree before and after insertion of "Adams"

B⁺-Tree Insertion
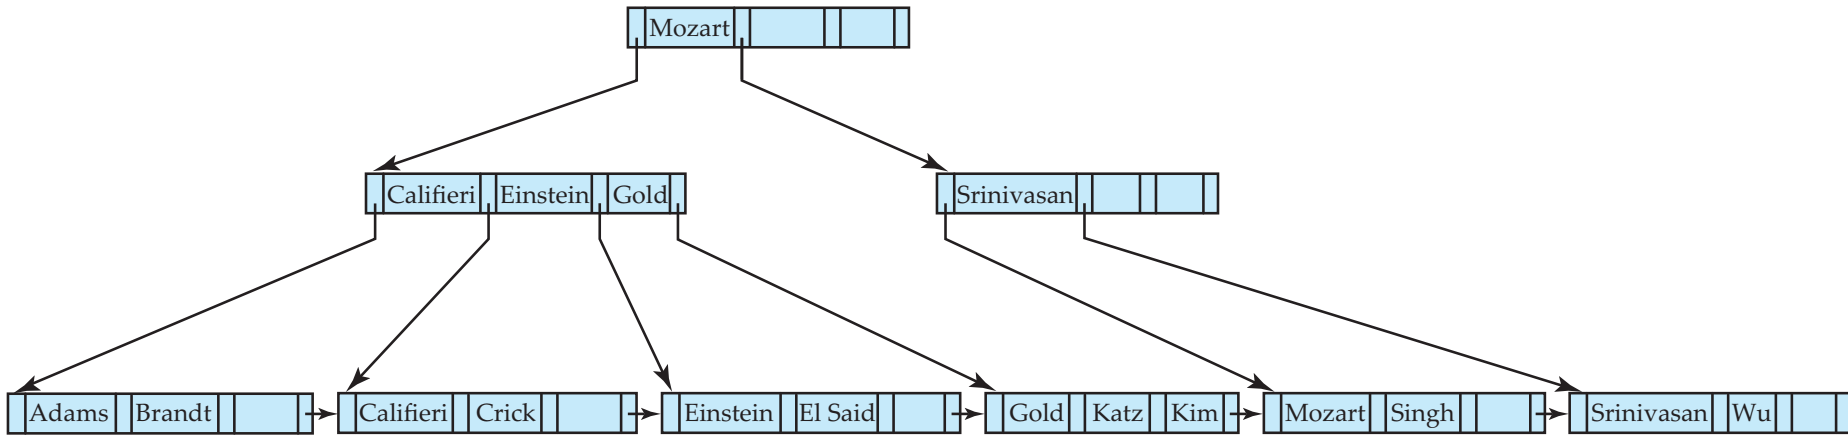
B⁺-Tree before and after insertion of "Lamport"

Insertion in B⁺-Trees (Cont.)

- Splitting a non-leaf node: when inserting (k,p) into an already full internal node N

    - Copy N to an in-memory area M with space for n+1 pointers and n keys

    - Insert (k,p) into M

    - Copy $P_1, K_1, \ldots, K_{\lceil n/2 \rceil-1}, P_{\lceil n/2 \rceil}$ from M back into node N

    - Copy $P_{\lceil n/2 \rceil+1}, K_{\lceil n/2 \rceil+1}, \ldots, K_n, P_{n+1}$ from M into newly allocated node N'

    - Insert $(K_{\lceil n/2 \rceil}, N')$ into parent N

- **Read pseudocode in book!**

Adams	Brandt	Califieri	Crick	

Califieri

Adams	Brandt	

|	Crick	

Examples of B⁺-Tree Deletion
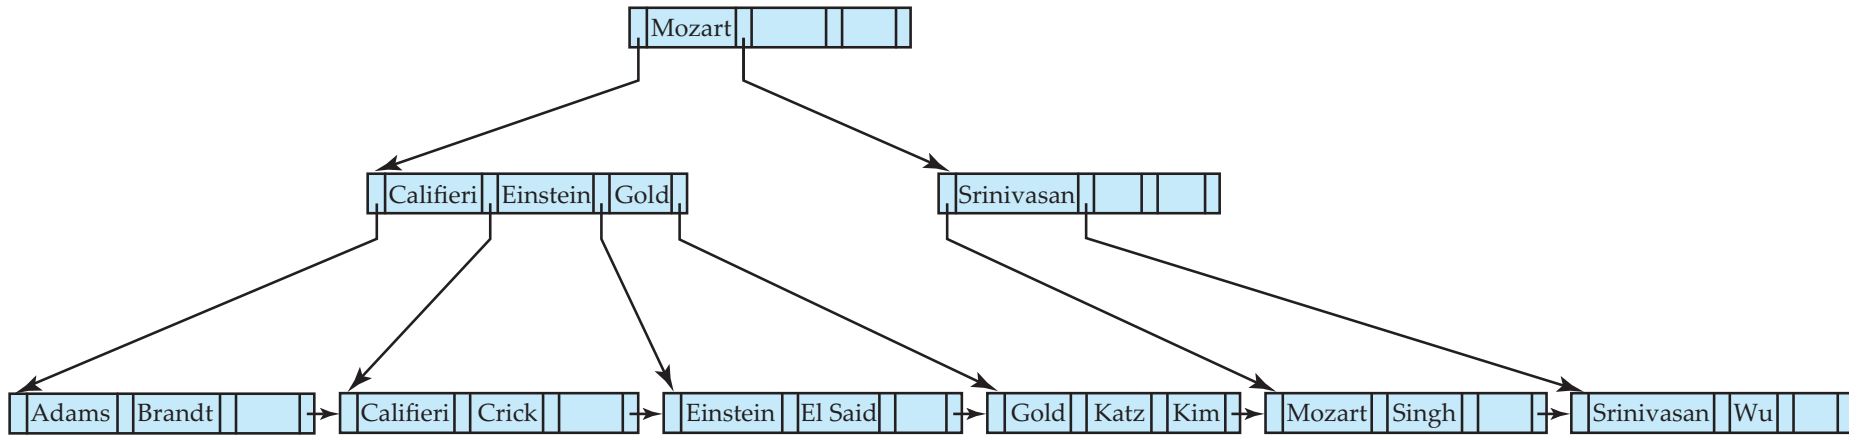
Before and after deleting "Srinivasan"

■ Deleting "Srinivasan" causes merging of under-full leaves

Examples of B+-Tree Deletion (Cont.)

Deletion of "Singh" and "Wu" from result of previous example

- Leaf containing Singh and Wu became underfull, and borrowed a value Kim from its left sibling
- Search-key value in the parent changes as a result

Example of B+-tree Deletion (Cont.)

Before and after deletion of "Gold" from earlier example

- Node with Gold and Katz became underfull, and was merged with its sibling
- Parent node becomes underfull, and is merged with its sibling
  - Value separating two nodes (at the parent) is pulled down when merging
- Root node then has only one child, and is deleted

Updates on B+-Trees: Deletion

■ Find the record to be deleted, and remove it from the main file and from the bucket (if present)

■ Remove (search-key value, pointer) from the leaf node if there is no bucket or if the bucket has become empty

■ If the node has too few entries due to the removal, and the entries in the node and a sibling fit into a single node, then **_merge siblings_**:

● Insert all the search-key values in the two nodes into a single node (the one on the left), and delete the other node.

● Delete the pair (K_{i-1}, P_i), where P_i is the pointer to the deleted node, from its parent, recursively using the above procedure.

Updates on B+-Trees: Deletion

- Otherwise, if the node has too few entries due to the removal, but the entries in the node and a sibling do not fit into a single node, then **redistribute pointers**:
  - Redistribute the pointers between the node and a sibling such that both have more than the minimum number of entries.
  - Update the corresponding search-key value in the parent of the node.

- The node deletions may cascade upwards till a node which has $\lceil n/2 \rceil$ or more pointers is found.

- If the root node has only one pointer after deletion, it is deleted and the sole child becomes the root.

Non-Unique Search Keys

- Alternatives to scheme described earlier
    - Buckets on separate block (bad idea)
    - List of tuple pointers with each key
        - ▸ Extra code to handle long lists
        - ▸ Deletion of a tuple can be expensive if there are many duplicates on search key (why?)
        - ▸ Low space overhead, no extra cost for queries
    - Make search key unique by adding a record-identifier
        - ▸ Extra storage overhead for keys
        - ▸ Simpler code for insertion/deletion
   - ▸ Widely used

B+-Tree File Organization

- Index file degradation problem is solved by using B+-Tree indices.

- Data file degradation problem is solved by using B+-Tree File Organization.

- The leaf nodes in a B+-tree file organization store records, instead of pointers.

- Leaf nodes are still required to be half full
  - Since records are larger than pointers, the maximum number of records that can be stored in a leaf node is less than the number of pointers in a nonleaf node.

- Insertion and deletion are handled in the same way as insertion and deletion of entries in a B+-tree index.

B⁺-Tree File Organization (Cont.)

Example of B⁺-tree File Organization

- Good space utilization important since records use more space than pointers.

- To improve space utilization, involve more sibling nodes in redistribution during splits and merges

  - Involving 2 siblings in redistribution (to avoid split / merge where possible) results in each node having at least $\lfloor 2n/3 \rfloor$ entries

Hashing

Static Hashing

■ A **bucket** is a unit of storage containing one or more records (a bucket is typically a disk block).

■ In a **hash file organization** we obtain the bucket of a record directly from its search-key value using a **hash function.**

■ Hash function h is a function from the set of all search-key values K to the set of all bucket addresses B.

■ Hash function is used to locate records for access, insertion as well as deletion.

■ Records with different search-key values may be mapped to the same bucket; thus entire bucket has to be searched sequentially to locate a record.

Example of Hash File Organization

Hash file organization of *instructor* file, using *dept_name* as key
(See figure in next slide.)

■ There are 10 buckets,

■ The binary representation of the ith character is assumed to be the integer i.

■ The hash function returns the sum of the binary representations of the characters modulo 10

  ● E.g. h(Music) = 1   h(History) = 2
        h(Physics) = 3   h(Elec. Eng.) = 3

Example of Hash File Organization

bucket 0

bucket 1

15151	Mozart	Music	40000

bucket 2

32343	El Said	History	80000
58583	Califieri	History	60000

bucket 3

22222	Einstein	Physics	95000
33456	Gold	Physics	87000
98345	Kim	Elec. Eng.	80000

bucket 4

12121	Wu	Finance	90000
76543	Singh	Finance	80000

bucket 5

76766	Crick	Biology	72000

bucket 6

10101	Srinivasan	Comp. Sci.	65000
45565	Katz	Comp. Sci.	75000
83821	Brandt	Comp. Sci.	92000

bucket 7

Hash file organization of *instructor* file, using *dept_name* as key (see previous slide for details).

Hash Functions

■ Worst hash function maps all search-key values to the same bucket; this makes access time proportional to the number of search-key values in the file.

■ An ideal hash function is **uniform**, i.e., each bucket is assigned the same number of search-key values from the set of *all* possible values.

■ Ideal hash function is **random**, so each bucket will have the same number of records assigned to it irrespective of the *actual distribution* of search-key values in the file.

■ Typical hash functions perform computation on the internal binary representation of the search-key.

● For example, for a string search-key, the binary representations of all the characters in the string could be added and the sum modulo the number of buckets could be returned. .

Handling of Bucket Overflows

- Bucket overflow can occur because of
  - Insufficient buckets
  - Skew in distribution of records. This can occur due to two reasons:
    - multiple records have same search-key value
    - chosen hash function produces non-uniform distribution of key values
- Although the probability of bucket overflow can be reduced, it cannot be eliminated; it is handled by using *overflow buckets*.

Handling of Bucket Overflows (Cont.)

- **Overflow chaining** – the overflow buckets of a given bucket are chained together in a linked list.

- Above scheme is called **closed hashing.**

  - An alternative, called **open hashing**, which does not use overflow buckets, is not suitable for database applications.


overflow buckets for bucket 1

Hash Indices

- Hashing can be used not only for file organization, but also for index-structure creation.

- A **hash index** organizes the search keys, with their associated record pointers, into a hash file structure.

- Strictly speaking, hash indices are always secondary indices
  - if the file itself is organized using hashing, a separate primary hash index on it using the same search-key is unnecessary.
  - However, we use the term hash index to refer to both secondary index structures and hash organized files.

Example of Hash Index
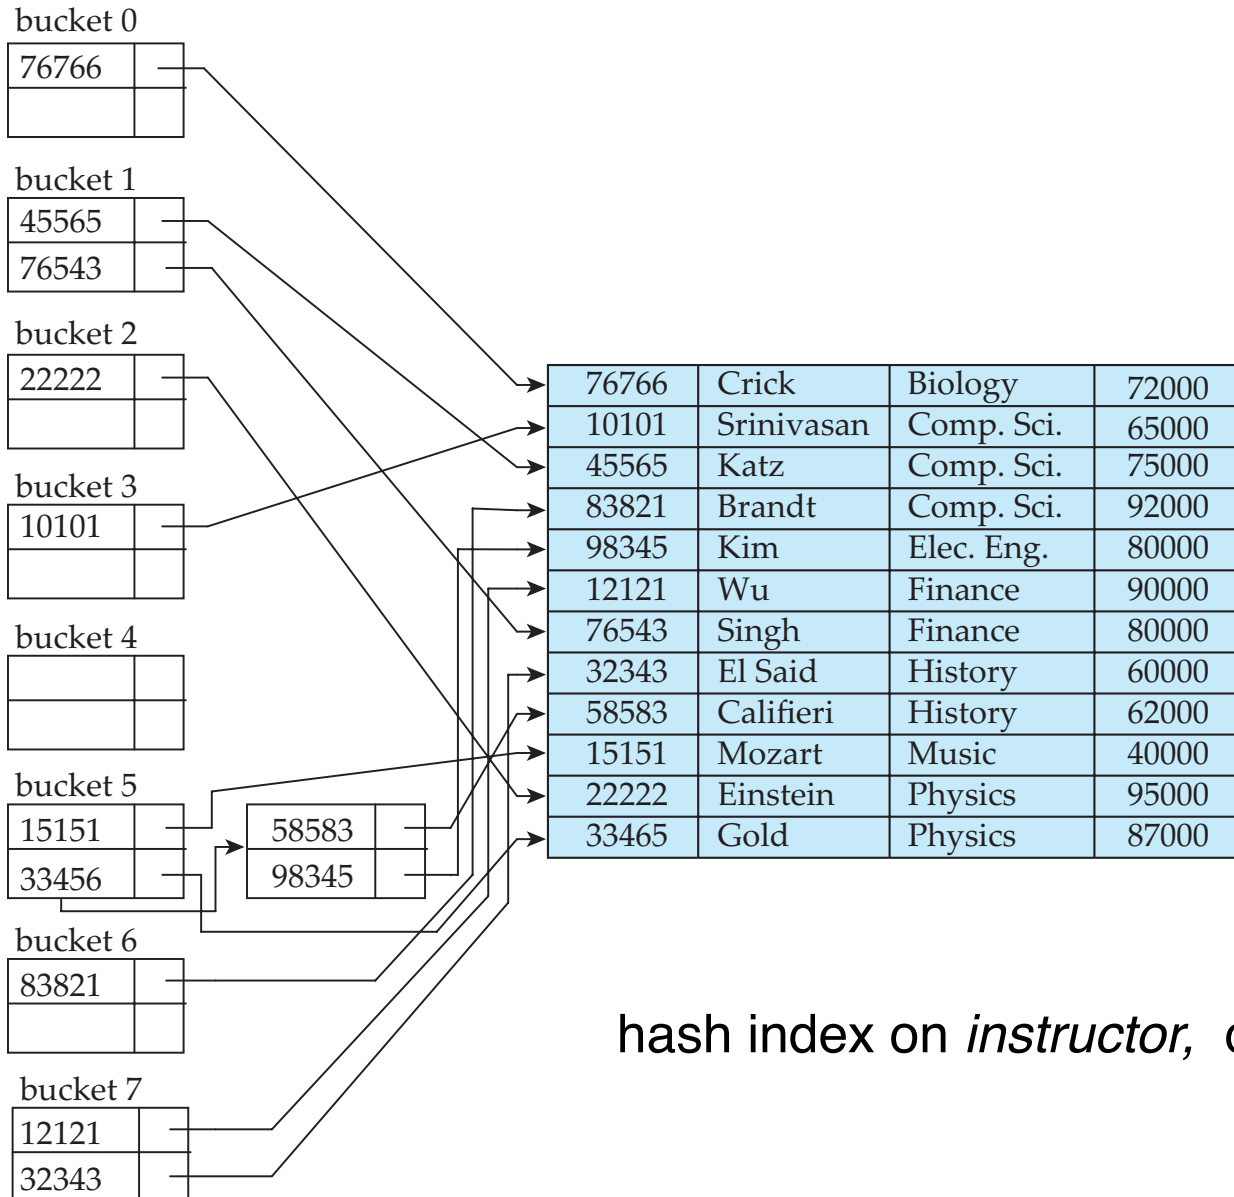
hash index on *instructor,* on attribute *ID*

Deficiencies of Static Hashing

- In static hashing, function *h* maps search-key values to a fixed set of *B* of bucket addresses. Databases grow or shrink with time.
  - If initial number of buckets is too small, and file grows, performance will degrade due to too much overflows.
  - If space is allocated for anticipated growth, a significant amount of space will be wasted initially (and buckets will be underfull).
  - If database shrinks, again space will be wasted.
- One solution: periodic re-organization of the file with a new hash function
  - Expensive, disrupts normal operations
- Better solution: allow the number of buckets to be modified dynamically.

Index Definition in SQL

- Create an index

    create index <index-name> **on** <relation-name>
    (<attribute-list>)

    E.g.: **create index** *b-index* **on** *branch(branch_name)*

- Use **create unique index** to indirectly specify and enforce the condition that the search key is a candidate key is a candidate key.

    - Not really required if SQL **unique** integrity constraint is supported

- To drop an index

    drop index <index-name>

- Most database systems allow specification of type of index, and clustering.

End of Chapter

Figure 11.01

10101	Srinivasan	Comp. Sci.	65000	
12121	Wu	Finance	90000	
15151	Mozart	Music	40000	
22222	Einstein	Physics	95000	
32343	El Said	History	60000	
33456	Gold	Physics	87000	
45565	Katz	Comp. Sci.	75000	
58583	Califieri	History	62000	
76543	Singh	Finance	80000	
76766	Crick	Biology	72000	
83821	Brandt	Comp. Sci.	92000	
98345	Kim	Elec. Eng.	80000	

Figure 11.15

Partitioned Hashing

- Hash values are split into segments that depend on each attribute of the search-key.

$$(A_1, A_2, \ldots, A_n) \text{ for } n \text{ attribute search-key}$$

- Example: $n = 2$, for *customer,* search-key being (*customer-street, customer-city*)

search-key value	hash value
(Main, Harrison)	101 111
(Main, Brooklyn)	101 001
(Park, Palo Alto)	010 010
(Spring, Brooklyn)	001 001
(Alma, Palo Alto)	110 010

- To answer equality query on single attribute, need to look up multiple buckets. Similar in effect to grid files.
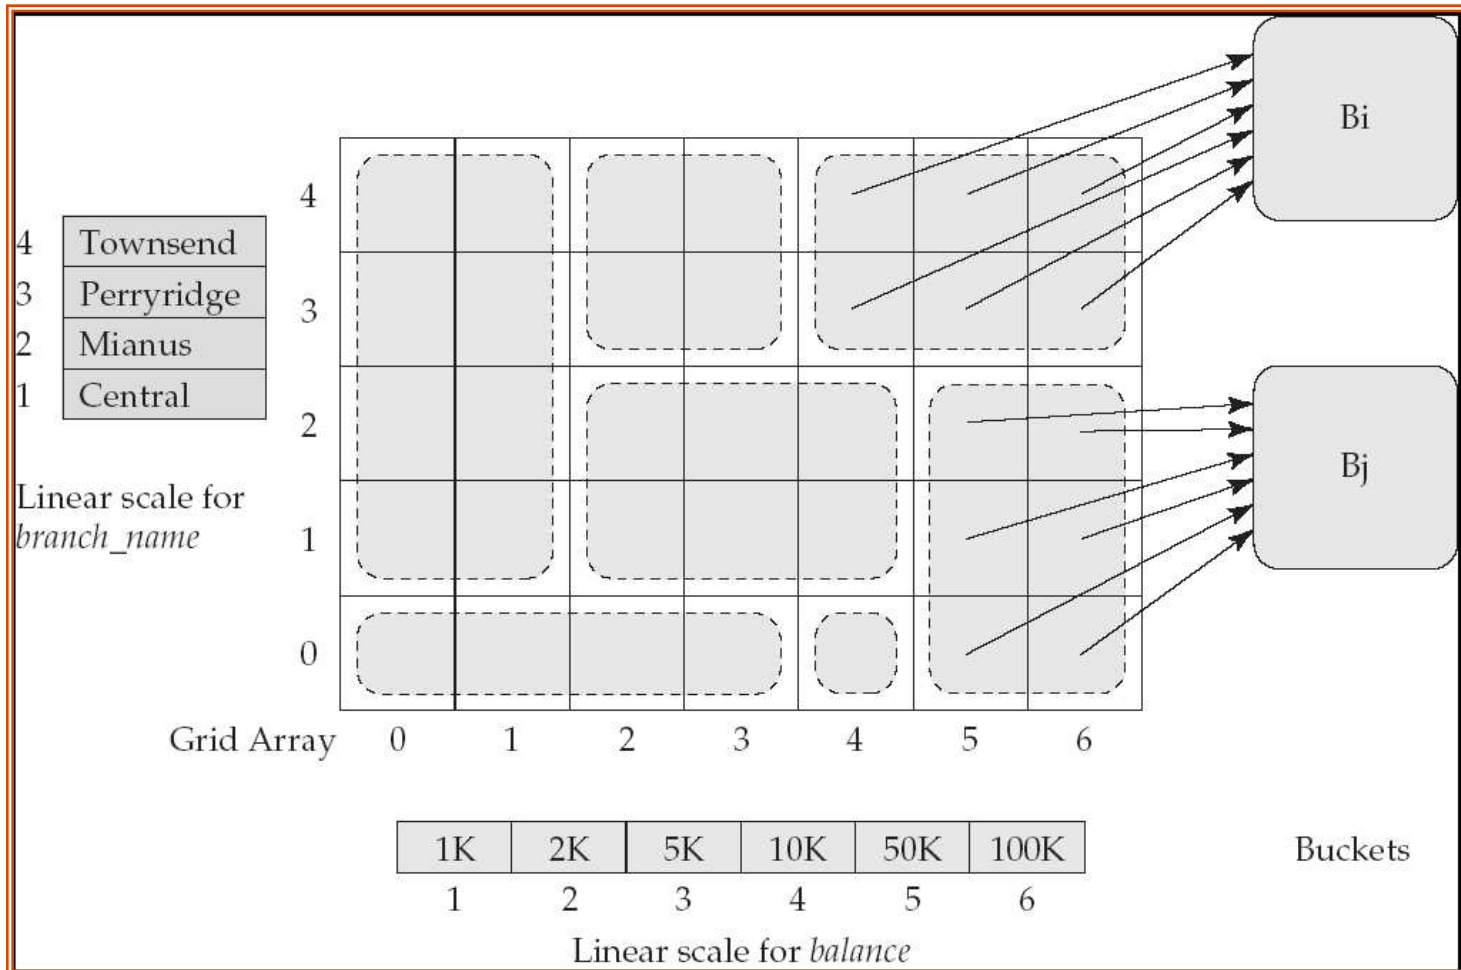
Grid Files

- Structure used to speed the processing of general multiple search-key queries involving one or more comparison operators.

- The grid file has a single grid array and one linear scale for each search-key attribute. The grid array has number of dimensions equal to number of search-key attributes.

- Multiple cells of grid array can point to same bucket

- To find the bucket for a search-key value, locate the row and column of its cell using the linear scales and follow pointer

Example Grid File for *account*

Queries on a Grid File

- A grid file on two attributes A and B can handle queries of all following forms with reasonable efficiency

  - $(a_1 \le A \le a_2)$
  - $(b_1 \le B \le b_2)$
  - $(a_1 \le A \le a_2 \ \wedge \ b_1 \le B \le b_2)$,.

- E.g., to answer $(a_1 \le A \le a_2 \ \wedge \ b_1 \le B \le b_2)$, use linear scales to find corresponding candidate grid array cells, and look up all the buckets pointed to from those cells.

Grid Files (Cont.)

- During insertion, if a bucket becomes full, new bucket can be created if more than one cell points to it.

  - Idea similar to extendable hashing, but on multiple dimensions

  - If only one cell points to it, either an overflow bucket must be created or the grid size must be increased

- Linear scales must be chosen to uniformly distribute records across cells.

  - Otherwise there will be too many overflow buckets.

- Periodic re-organization to increase grid size will help.

  - But reorganization can be very expensive.

- Space overhead of grid array can be high.

- R-trees (Chapter 23) are an alternative